Stochastic Satisfiability Modulo Theory: A Novel Technique for the Analysis of Probabilistic Hybrid Systems [1]

Martin Fränzle, Tino Teige[2]

Carl von Ossietzky Universität, Oldenburg, Germany

Holger Hermanns[3]

Saarland University, Saarbrücken, Germany

Abstract

The analysis of hybrid systems exhibiting probabilistic behaviour is notoriously difficult. To enable mechanised analysis of such systems, we extend the reasoning power of arithmetic satisfiability-modulo-theory solving (SMT) by a comprehensive treatment of randomized (a.k.a. stochastic) quantification over discrete variables within the mixed Boolean-arithmetic constraint system. This provides the technological basis for a fully symbolic analysis of probabilistic hybrid automata. Generalizing SMT-based bounded model-checking of hybrid automata [2,9], stochastic SMT permits the direct and fully symbolic analysis of probabilistic bounded reachability problems of probabilistic hybrid automata without resorting to approximation by intermediate finite-state abstractions.

Keywords: Stochastic satisfiability, infinite domains, probabilistic hybrid automata, probabilistic bounded reachability.

Over the last decade, formal verification of digital systems has evolved from an academic subject to an approach accepted by industry, with dozens of commercial tools now available. Among the most successful verification methods for finite-state systems is *bounded model checking* (BMC), as suggested by Groote et al. in [11] and by Biere et al. in [3]. The idea of BMC is to encode the next-state relation of a system as a propositional formula, to unroll this to some given finite depth k, and to augment it with a corresponding finite unravelling of the tableau of (the negation of) a temporal formula in order to obtain a propositional SAT problem which is satisfiable if and only if an error trace of length k exists. Enabled by the impressive

[1] This work has been partially supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center "Automatic Verification and Analysis of Complex Systems" (SFB/TR 14 AVACS, www.avacs.org).

[2] Email: fraenzle@informatik.uni-oldenburg.de, teige@informatik.uni-oldenburg.de

[3] Email: hermanns@cs.uni-sb.de

This paper is electronically published in
Electronic Notes in Theoretical Computer Science
URL: www.elsevier.nl/locate/entcs

gains in performance of propositional SAT checkers in recent years, BMC can now be applied to very large finite-state designs.

Though originally formulated for discrete transition systems, the concept of BMC also applies to hybrid discrete-continuous systems. The BMC formulae arising from such systems comprise complex Boolean combinations of arithmetic constraints over real-valued variables, thus entailing the need for satisfiability-modulo-theory (SMT) solvers over arithmetic theories to solve them. Such SMT procedures are thus currently in the focus of the SAT-solving community (e.g., [8]), as is their application to and tailoring for BMC of hybrid systems (e.g., [2,9]).

The scope of these procedures, however, is confined to purely Boolean queries of the form "can the system ever exhibit an undesirable behavior?", whereas requirements for safety-critical systems frequently take the form of bounds on error probability, requiring the residual probability of engaging into undesirable behavior to be below an acceptable threshold. Automatically answering such queries requires, first, models of hybrid behavior that are able to represent probabilistic effects like component breakdown and, second, algorithms for state space traversal of such hybrid models.

In the context of hybrid systems augmented with probabilities, a wealth of models has been suggested by various authors. These models vary with respect to the degree of continuous dynamics, the support for random phenomena, and the degree to which they support non-determinism and compositionality. The cornerstones are formed by *probabilistic hybrid automata*, where state changes forced by continuous dynamics may involve discrete random experiments [15], *piecewise deterministic Markov processes* [7], where state changes may happen spontaneously in a manner similar to continuous-time Markov processes, and *stochastic differential equations* [1], where, like in Brownian motion, the random perturbation affects the dynamics continuously. In full generality, stochastic hybrid system (SHS) models can cover all such ingredients [6]. While such models have a vast potential of application, results related to their analysis and verification are limited, and often based on Monte-Carlo simulation [4,12]. For certain subclasses of piecewise deterministic Markov processes, of probabilistic hybrid automata, and of stochastic hybrid systems, reachability probabilities can be approximated (e.g. [15,5,13]).

In [10], we presented a technology that saves the virtues of SMT-based BMC, namely the fully symbolic treatment of hybrid state spaces, while advancing the reasoning power to probabilistic models and requirements. While the technique is more general, the current paper focuses on depth-bounded reachability of discrete-time probabilistic hybrid automata. With respect to the stochastic dynamics considered this model is very simple and thus constitutes a good attack point to pioneer effective model checking techniques for probabilistic hybrid systems, harvesting recent advances in depth-bounded reachability analysis for ordinary hybrid systems. Albeit being simple, the model of probabilistic hybrid automata has interesting practical applications [15].

In order to achieve this, in [10] we introduced *stochastic satisfiability modulo theory* (SSMT) as the unification of stochastic propositional satisfiability (SSAT) [14] and satisfiability modulo theory. The SSMT framework deals with *existential* and *randomized* quantification of finite-domain variables. An SSMT formula is specified

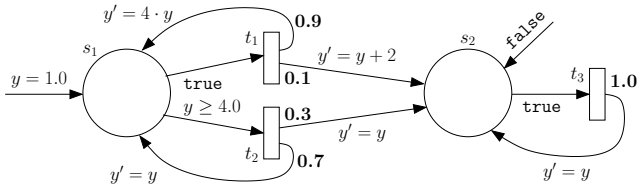

Fig. 1. A probabilistic hybrid automaton $\mathcal{H}$.

by a quantifier prefix and an SMT formula, e.g. $\Phi = \exists x \in \{0,1\}\ \mathfrak{U}_{\langle(0,0.6),(1,0.4)\rangle} y \in \{0,1\} : (x > 0 \vee 2a + 4b \geq 3) \wedge (y > 0 \vee 2a + 4b < 1)$. The value of a variable bound by an existential quantifier, as in $\exists x \in \{0,1\}$, can be set arbitrarily, while the value of a variable bound by a randomized quantifier, as in $\mathfrak{U}_{\langle(0,0.6),(1,0.4)\rangle} y \in \{0,1\}$, is determined stochastically by the corresponding distribution, here $\langle(0,0.6),(1,0.4)\rangle$. For instance, $\mathfrak{U}_{\langle(0,0.6),(1,0.4)\rangle} y \in \{0,1\}$ means that the variable y is assigned the value 0 or 1 with probability 0.6 or 0.4, respectively. The solution of an SSMT problem Φ is a strategy to assign values to the existential variables that *maximizes the overall satisfaction probability* of Φ. Since the quantifier prefix of Φ allows an alternating sequence of existential and randomized quantifiers, the value of an existential variable depends on the values of the randomized variables with earlier appearance in the prefix. Consequently, in general such a solution is a tree of assignments to the existential variables depending on the values of preceding randomized variables. For the SSMT formula $\Phi = \exists x \in \{0,1\}\ \mathfrak{U}_{\langle(0,0.6),(1,0.4)\rangle} y \in \{0,1\} : \varphi$, the goal is to determine the maximum probability s.t. there is a value for x s.t. for random values of y the SMT formula φ is satisfiable. More formally, the *maximum probability of satisfaction* $Pr(\Phi)$ of an SSMT formula Φ is defined recursively as follows, where φ denotes the SMT formula.

1. $Pr(\varphi) = 0$ if φ is unsatisfiable, and 1 otherwise.
2. $Pr(\exists x_i \in \mathrm{dom}(x_i) \ldots Q_n x_n \in \mathrm{dom}(x_n) : \varphi)$
   $= \max_{v \in \mathrm{dom}(x_i)} Pr(Q_{i+1} x_{i+1} \in \mathrm{dom}(x_{i+1}) \ldots Q_n x_n \in \mathrm{dom}(x_n) : \varphi[v/x_i])$.
3. $Pr(\mathfrak{U}_{d_i} x_i \in \mathrm{dom}(x_i) \ldots Q_n x_n \in \mathrm{dom}(x_n) : \varphi)$
   $= \sum_{(v,p) \in d_i} p \cdot Pr(Q_{i+1} x_{i+1} \in \mathrm{dom}(x_{i+1}) \ldots Q_n x_n \in \mathrm{dom}(x_n) : \varphi[v/x_i])$.

In [10], we proposed an algorithm for solving SSMT problems in the sense of determining the maximum probability of satisfaction. This algorithm extends solvers for SSAT [14] in much the same way that DPLL(T) solvers (e.g., [8]) extend classical DPLL SAT solvers. First experimental results prove the concept of our approach and show the impact of algorithmic acceleration techniques for SSMT problems.

A *discrete-time probabilistic hybrid automaton* (PHA) as described, e.g., in [10] extends the notion of a hybrid automaton, where the non-deterministic selection of a transition is enriched by a probabilistic choice according to a distribution over variants of the transition. I.e., each transition carries a (discrete) probabilistic distribution. Each probabilistic choice within such a distribution leads to a potentially different successor mode while performing some discrete actions, cf. Fig. 1 for an example. We are especially interested in k-bounded model checking problems, i.e., we want to prove or disprove whether a given reachability property P is satisfied with a maximum probability greater or equal p in a probabilistic hybrid automaton

along all its traces of length up to k. For an illustration of *probabilistic bounded reachability* consider the probabilistic hybrid automaton $\mathcal{H}$ from Fig. 1 where mode s_2 is the reachability goal. The maximum probabilities of reaching s_2 in 0, 1, 2, and 3 steps are 0.0, 0.1, $0.1 + 0.9 \cdot 0.3 = 0.37$, and $0.1 + 0.9 \cdot (0.3 + 0.7 \cdot 0.3) = 0.559$, respectively.

The idea of the formalized encoding of a PHA $\mathcal{H}$ into an SSMT formula Φ, as presented in [10], is that the non-deterministic choice of a transition in a PHA corresponds to existential quantification in SSMT, while the probabilistic distributions correspond to randomized quantification. The discrete-continuous behavior of the automaton then is encoded by means of standard techniques. The construction of Φ ensures that Φ is satisfiable with maximum probability p iff the PHA $\mathcal{H}$ (restricted to traces of length k) fulfills a certain property P with maximum probability p. Hence, we can reduce the probabilistic bounded reachability problem of PHAs to the SSMT problem.

This symbolic encoding together with the Stochastic SMT procedure provides fully symbolic analysis of probabilistic bounded reachability problems of probabilistic hybrid automata without resorting to approximation by intermediate finite-state abstractions.

References

[1] L. Arnold. *Stochastic Differential Equations: Theory and Applications*. Wiley - Interscience, 1974.

[2] G. Audemard, M. Bozzano, A. Cimatti, and R. Sebastiani. Verifying industrial hybrid systems with MathSAT. In *Bounded Model Checking (BMC'04)*, volume 119 of *ENTCS*, pages 17–32, 2004.

[3] A. Biere, A. Cimatti, and Y. Zhu. Symbolic model checking without BDDs. In *TACAS'99*, volume 1579 of *LNCS*. Springer Verlag, 1999.

[4] H. A. P. Blom, J. Krystul, and G. J. Bakker. A particle system for safety verification of free flight in air traffic. In *Decision and Control*, pages 1574–1579. IEEE, 2006.

[5] M. L. Bujorianu and J. Lygeros. Reachability questions in piecewise deterministic markov processes. In *Hybrid Systems: Computation and Control*, volume 2623 of *LNCS*, pages 126–140. Springer, 2003.

[6] M. L. Bujorianu and J. Lygeros. Toward a general theory of stochastic hybrid systems. In *Stochastic Hybrid Systems: Theory and Safety Critical Applications*, volume 337 of *LNCIS*, pages 3–30. Springer, 2006.

[7] M. Davis. *Markov Models and Optimization*. Chapman & Hall, London, 1993.

[8] B. Dutertre and L. de Moura. A Fast Linear-Arithmetic Solver for DPLL(T). In *Computer-Aided Verification*, volume 4144 of *LNCS*, pages 81–94. Springer, 2006.

[9] M. Fränzle and C. Herde. HySAT: An efficient proof engine for bounded model checking of hybrid systems. *Formal Methods in System Design*, 30:179–198, 2007.

[10] M. Fränzle, H. Hermanns, and T. Teige. Stochastic Satisfiability Modulo Theory: A Novel Technique for the Analysis of Probabilistic Hybrid Systems. In *Proceedings of the 11th International Conference on Hybrid Systems: Computation and Control (HSCC'08)*, 2008. to appear.

[11] J. F. Groote, J. W. C. Koorn, and S. F. M. van Vlijmen. The Safety Guaranteeing System at Station Hoorn-Kersenboogerd. In *Conference on Computer Assurance*, pages 57–68. National Institute of Standards and Technology, 1995.

[12] J. P. Hespanha. Polynomial stochastic hybrid systems. In *Hybrid Systems: Computation and Control*, volume 3414 of *LNCS*, pages 322–338. Springer, 2005.

[13] X. D. Koutsoukos and D. Riley. Computational methods for reachability analysis of stochastic hybrid systems. In J. P. Hespanha and A. Tiwari, editors, *HSCC*, volume 3927 of *LNCS*, pages 377–391. Springer Verlag, 2006.

[14] C. H. Papadimitriou. Games against nature. *J. Comput. Syst. Sci.*, 31(2):288–301, 1985.

[15] J. Sproston. *Model Checking of Probabilistic Timed and Hybrid Systems*. PhD thesis, University of Birmingham, 2000.