SBIP 2.0: Statistical Model Checking Stochastic Real-time Systems

Braham Lotfi Mediouni, Ayoub Nouri, Marius Bozga, Mahieddine Dellabani,
Axel Legay, Saddek Bensalem

► To cite this version:

Braham Lotfi Mediouni, Ayoub Nouri, Marius Bozga, Mahieddine Dellabani, Axel Legay, et al.. SBIP 2.0: Statistical Model Checking Stochastic Real-time Systems. ATVA 2018 - 16th International Symposium Automated Technology for Verification and Analysis, Oct 2018, Los Angeles, CA, United States. pp.536-542, 10.1007/978-3-030-01090-4_33 . hal-01888538

HAL Id: hal-01888538

<https://hal.science/hal-01888538>

Submitted on 5 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

SBIP 2.0: Statistical Model Checking Stochastic Real-time Systems^{*}

Braham Lotfi Mediouni¹, Ayoub Nouri¹, Marius Bozga¹, Mahieddine Dellabani¹, Axel Legay², and Saddek Bensalem¹

¹ Univ. Grenoble Alpes, CNRS, Grenoble INP^{**}, VERIMAG, 38000 Grenoble, France

² INRIA, Rennes, France

Abstract. This paper presents a major new release of SBIP, an extensible statistical model checker for Metric (MTL) and Linear-time Temporal Logic (LTL) properties on respectively Generalized Semi-Markov Processes (GSMP), Continuous-Time (CTMC) and Discrete-Time Markov Chain (DTMC) models. The newly added support for MTL, GSMPs, CTMCs and rare events allows to capture both real-time and stochastic aspects, allowing faithful specification, modeling and analysis of real-life systems. SBIP is redesigned as an IDE providing project management, model edition, compilation, simulation, and statistical analysis.

1 Introduction

Statistical Model Checking (SMC) is a powerful alternative to classical numerical probabilistic model-checking that generally fail to handle large state-space systems. SMC was successfully applied in the assessment of different real-life systems in various application domains. Classical model checkers [8, 4] now include SMC as part of their analysis engines and have been recently joined by a variety of specialized ones [12, 6, 1, 9]. All these tools mainly differ in their modeling and properties specification formalisms. UPPAAL-SMC [4] considers *Networks of Priced Timed Automata*, which are high-level representations of D/CTMCs for system modeling, and weighted MTL for properties specification. PRISM [8] treats in addition *Markov Decision Processes* and *Probabilistic Timed Automata* for modeling, and *Probabilistic Computation Tree*, *Continuous Stochastic Logic* (CSL), and LTL for specification. Plasma Lab [6] is a modular statistical model checker that allows to use external simulators and checkers. Its default configuration supports DTMCs specified in a PRISM dialect and bounded LTL. Ymer [12] is one of the rare tools to implement SMC (Hypothesis testing) for GSMPs and CSL, however it is no more maintained. Finally, COSMOS [1] relies on *Generalized Stochastic Petri Nets* as input models and *Hybrid Automata Stochastic Logic*, a more expressive formalism, for properties specifications.

In this paper, we present the newest release of SBIP, a statistical model checker that enriches the existing BIP tool-set [2] with statistical analyses. BIP

^{*} The research leading to these results has received funding from the EU's H2020 programme under grant agreements no. 700665 (CITADEL), 7300080 (ESROCOS)

^{**} Institute of Engineering Univ. Grenoble Alpes

provides a general framework to support design activities ranging from specification and validation to implementation and deployment in a rigorous way. To implement this vision, a rich tool-set was built for modeling, languages embedding, functional validation, models transformation and distributed code generation.

In its previous version [9], **SBIP** was limited to the analysis of DTMCs with respect to bounded LTL properties. In this release, it was redesigned and extended to support **GSMPs, CTMCs, MTL, parametric exploration of LTL and MTL properties and analysis of rare events**. The tool has also benefited from a major revision of its workflows and GUI. It now provides an Integrated Development Environment (IDE) where one can edit, compile, simulate models, and perform analyses. Additionally, **SBIP** is now organized around well-structured projects that enclose models, properties and traces. It also includes support for graphical visualization of analysis results.

2 **SBIP** Design and Functionalities

SBIP is fully developed in Java and runs on GNU/Linux. It is freely available at <http://www-verimag.imag.fr/Statistical-Model-Checking.html>. The tool is distributed with a large set of case studies and a detailed documentation (e.g., user manual, installation details, video tutorials). For the sake of simplicity, we also provide a virtual machine with a pre-installed version of the tool.

This new release was designed in a modular fashion to allow more flexibility and extensibility. As depicted in Fig. 1, **SBIP** consists of three generic functional modules: *Stochastic Simulation Engine*, *Monitoring*, and *Statistical Analyses* that currently include *Hypothesis Testing (HT)*, *Probability Estimation (PE)*, *Parametric Exploration (PX)* and *Importance Splitting (IP)* for rare events analysis. All these modules are fully independent and interact through well-defined Java interfaces. The latter also define a clean and easy way to extend the tool with further modules (simulators, monitors and analyzers). In practice, statistical analysis algorithms trigger the stochastic simulation engine to produce a new execution trace which is monitored against an input property to produce a local verdict. Depending on the used analysis method, several iterations are generally required, to produce the final verdict. The proposed design allows to perform different analyses in separate workflows, namely simple simulation, standard SMC analyses, parametric SMC exploration and analysis of rare events. These workflows rely on common features such as models and properties edition, compilation and generated traces inspection.

Stochastic Simulation Engine. Currently, **SBIP** allows to use two different stochastic simulators, namely, for classical stochastic BIP [9] that enables to model discrete-time systems (DTMCs) and for the newly implemented Stochastic Real-Time BIP [10] for continuous-time systems with arbitrary distributions (GSMPs and CTMCs)³. The former produces untimed traces needed to verify bounded LTL properties (and to guarantee backward compatibility), whereas

³ SRT-BIP sources are available at <https://gricad-gitlab.univ-grenoble-alpes.fr/verimag/bip/compiler/tree/stochastic-real-time>

Fig. 1: SBIP architecture

the latter generates timed traces necessary to verify MTL properties. We implemented simulators to produce traces in different modes, i.e., symbol-wise, piece-wise and trace-wise. We use the first mode for online monitoring and to be able to interrupt simulations as soon as a verdict is obtained. The second is primordial for rare events analysis and allows to generate traces as a concatenation of trace-fragments. Finally, we use the third mode for offline monitoring.

Monitor. The new release of the tool implements monitoring capabilities for MTL and bounded LTL formulas. Our monitoring algorithms are inspired from the rewrite-based procedures introduced in [3, 11]. Given a formula and a trace, the monitor alternates rewriting and simplification phases. Rewriting consumes a symbol of the trace and partially evaluates the formula by unfolding temporal operators and evaluating atomic propositions to their truth value. Simplification applies Boolean reduction rules to the formula in order to conclude or to simplify it. The implemented MTL/LTL grammars and monitors allow for expressing properties with nested operators and having parameters, i.e., variables used to represent a range of properties in a compact way.

Statistical Analyses. In addition to classical SMC algorithms, i.e., *HT* [12] and *PE* [5], we propose in this release two additional analyses (exploitable via independent workflows) for the exploration of properties parameters, *Parametric Exploration (PX)*, and for rare events analysis, *Importance Splitting (IP)* [7]. To recall, *HT* allows to answer qualitative queries, i.e., given a stochastic system S and a property ϕ , it enables to assess whether the probability for S to satisfy ϕ is greater or equal to a given threshold θ . *PE* addresses quantitative queries, that is to compute a probability estimate p for S to satisfy ϕ .

Parametric Exploration (PX) is an automated way to perform statistical model checking on a family of properties, in a batch mode. A family of properties is specified in a compact way as a parametric property $\phi(x)$, where x is an integer parameter ranging over a finite instantiation domain Π . Similarly to PRISM, our implemented algorithm returns a set of SMC verdicts corresponding to the verification of the parametric property instances $\phi(v_x)$ with respect to $v_x \in \Pi$.

This can be very useful when exploring unknown system parameters such as, buffers sizes guaranteeing no overflow, or the amount of consumed energy. It automates the exploration for large parameters domains as opposed to tedious and time consuming manual procedures. This exploration differs from UPPAAL-SMC parametric SMC which explores the parameters of the input model.

Importance Splitting (IP) overcomes the problem of estimating the probability $P(S \models \phi)$ of a system S to satisfy a property ϕ representing a rare event. This is done by considering a set of intermediate levels l_i that corresponds to less rare properties ϕ_i , s.t., $\phi_n \Rightarrow \phi_{n-1} \Rightarrow \dots \Rightarrow \phi_1$, where $\phi_n = \phi$. $P(S \models \phi)$ is therefore computed as the product of the conditional probabilities to reach l_i from l_{i-1} , i.e., $\prod_{i=1}^n P(S \models \phi_i \mid S \models \phi_{i-1})$. In our implementation, the intermediate levels l_i and associated ϕ_i are defined via a score function given as input. To evaluate a system trace with respect to ϕ , we implemented a procedure that tells the level reached by the trace, i.e., the intermediate property it satisfies. Our algorithm is similar to the analysis procedure proposed in Plasma Lab. It iterates over levels, and for each one, it simulates m trace prefixes among which m_s reach the next level and m_f do not. The conditional probability to reach the next level is thus estimated as the ratio m_s/m . In the next iteration, the simulation of successful prefixes is resumed, while the rest (m_f) are replaced by successful ones sampled uniformly. We note that *IP* is currently limited to the analysis of DTMCs.

3 Case Studies

In this section, we briefly present experiments performed using *SBIP*⁴. Different case studies covering various application domains were considered to validate the new release of the tool. We implemented models for communication protocols, namely Firewire, Bluetooth, and the Precision Time Protocol (PTP), for a vehicle gear controller, a Pacemaker and a mutual exclusion scenario. All the experiments were performed on a Dell Latitude 5480 with an i7-7820HQ processor and 32 GB of RAM, running Ubuntu 16.04.

On these models, we tackled different types of requirements. For the Firewire case study, we focused on analyzing its leader election protocol in different topologies (2, 3 and 5 nodes) with respect to convergence time, by considering the impact of contention ($\phi_{1,2,3}$) and regarding the impact of a node position on its probability to become the leader (ϕ_4). In this study, except ϕ_3 performed using *PE*, the other properties were performed using *PX*. We also built a parametric model of the Bluetooth device discovery mechanism with one sender and one receiver that can be either in an active (v1) or a sniff mode (v2). For this model, we were interested in studying the energy consumption of the receiver in both modes (ϕ_6) in addition to the convergence time (ϕ_5). The PTP protocol was subject to the analysis of the maximal drift between the master and the slave clocks (ϕ_7).

⁴ See details in <http://www-verimag.imag.fr/TR/TR-2018-5.pdf>

For the gearbox system, we investigated the minimum and maximum time required to complete a gear change (ϕ_8). We also verified requirements regarding the time relationships between atrial and ventricular events in the pacemaker model ($\phi_{9,10}$). Analyses of the Bluetooth, PTP and the gearbox models were performed using *PX*, while we used *PE* for the Pacemaker. We also considered a model of three concurrent processes arbitrarily requesting access to a shared resource. In this case study, the goal was to estimate the probability that each process is able to access the resource 10 times within 30 system steps (rare property ϕ_{11}). Using our *IP* implementation, we obtained 2.35×10^{-7} in less than 13s, while it was not possible to observe the rare event using *PE* upon 3 minutes of execution.

In addition to these experiments summarized in Table 1, we report in the last two columns some performance measures of the tool, namely, the number of SMC loops performed for parametric exploration, and the average SMC time for a single loop. We observed that depending on the model size and the property complexity, the time varies from some seconds to a dozen of minutes, except for the pacemaker model where it took more than an hour. In this particular case, *PE* required 4883 long execution traces, representing approximately 8 minutes of real system execution.

4 Discussion

Most SMC tools [8, 4, 12, 6, 1] use dedicated abstract models as input for verification. In contrast, SBIP uses BIP, a full-fledged expressive component-based framework developed to support system design from specification to analysis and implementation. It allows for incrementally building complex systems from elementary components and offers real-time capabilities, in addition to high-level coordination and synchronization primitives e.g. multi-party interactions and priorities. Furthermore, it enables including external C++ code, e.g. for modeling complex data structures and integrating legacy code.

Case study	Model	ϕ	Analysis	#smc loops	avg smc time
Firewire(2)	CTMC	ϕ_1	<i>PX</i>	11	1m 21s
		ϕ_2	<i>PX</i>	9	1m 59s
		ϕ_3	<i>PE</i>	-	2m 28s
		ϕ_4	<i>PX</i>	2	3m 27s
Firewire(3)	CTMC	ϕ_1	<i>PX</i>	17	1m 53s
		ϕ_2	<i>PX</i>	11	3m 34s
		ϕ_3	<i>PE</i>	-	3m 38s
		ϕ_4	<i>PX</i>	3	4m 43s
Firewire(5)	CTMC	ϕ_1	<i>PX</i>	18	3m 54s
		ϕ_2	<i>PX</i>	17	12m 36s
		ϕ_3	<i>PE</i>	-	7m 23s
		ϕ_4	<i>PX</i>	5	10m 16s
Bluetooth v1	CTMC	ϕ_5	<i>PX</i>	9	2m 27s
		ϕ_6	<i>PX</i>	16	3m 11s
Bluetooth v2	CTMC	ϕ_5	<i>PX</i>	11	3m 0s
		ϕ_6	<i>PX</i>	14	13m 05s
PTP	GSMP	ϕ_7	<i>PX</i>	15	8m 42s
Gear Control	CTMC	ϕ_8	<i>PX</i>	11	54s
Pacemaker	CTMC	ϕ_9	<i>PE</i>	-	1h 28m
		ϕ_{10}	<i>PE</i>	-	1h 30m
Mutual Exclusion	DTMC	ϕ_{11}	<i>IP</i>	-	13s
			<i>PE</i>	-	3m 37s

Table 1: Summary of performance

We briefly discuss **SBIP** capabilities with respect to major SMC tools. Regarding the analyses, **SBIP** implements the *HT* and *PE* algorithms similarly to UPPAAL-SMC [4], PRISM [8] and Plasma Lab [6]. Besides, only PRISM offers a parametric functionality similar to *PX*. Furthermore, to the best of our knowledge only Plasma Lab and COSMOS [1] support rare events analysis. The former is the only one implementing *IP* as in our tool, while the latter rather relies on importance sampling. Our underlying modeling formalism allows for expressing arbitrary probability distributions over time. It offers built-in standard distributions, e.g. Normal, and a simple mechanism for specifying custom distributions. In contrast, PRISM is restricted to uniform and exponential distributions, whereas in UPPAAL-SMC one need to define such distributions manually by using a subset of the C language. The expressiveness of BIP together with the reliance on concrete executions result in lower runtime performance compared to UPPAAL-SMC and PRISM. Comparatively, the authors of Plasma Lab chose to focus on modularity at the expense of performance. In the future, we plan to optimize our simulation engine to improve the overall performance.

References

1. P. Ballarini, B. Barbot, M. Duflot, S. Haddad, and N. Pekergin. Hasl: A new approach for performance evaluation and model checking from concepts to experimentation. *Performance Evaluation*, 90:53 – 77, 2015.
2. A. Basu, B. Bensalem, M. Bozga, J. Combaz, M. Jaber, T. H. Nguyen, and J. Sifakis. Rigorous component-based system design using the bip framework. *IEEE Software*, 28(3):41–48, May 2011.
3. P. E. Bulychiev, A. David, K. G. Larsen, A. Legay, G. Li, and D. B. Poulsen. Rewrite-based statistical model checking of wmtl. *RV*, 7687:260–275, 2012.
4. A. David, K. G. Larsen, A. Legay, M. Mikušionis, and D. B. Poulsen. Uppaal smc tutorial. *STTT*, 17(4):397–415, August 2015.
5. T. Hérault, R. Lassaigne, F. Magniette, and S. Peyronnet. Approximate Probabilistic Model Checking. In *VMCAI’04*, pages 73–84, January 2004.
6. C. Jegourel, A. Legay, and S. Sedwards. A platform for high performance statistical model checking – plasma. In *TACAS’12*, pages 498–503, Berlin, 2012. Springer.
7. C. Jegourel, A. Legay, and S. Sedwards. Importance splitting for statistical model checking rare properties. In *CAV*, volume 13, pages 576–591. Springer, 2013.
8. M. Kwiatkowska, G. Norman, and D. Parker. Prism 4.0: verification of probabilistic real-time systems. *CAV’11*, pages 585–591. Springer-Verlag, 2011.
9. A. Nouri, S. Bensalem, M. Bozga, B. Delahaye, C. Jegourel, and A. Legay. Statistical model checking QoS properties of systems with SBIP. *STTT’15*, 17:171–185.
10. A. Nouri, B. L. Mediouni, M. Bozga, J. Combaz, A. Legay, and S. Bensalem. Performance evaluation of stochastic real-time systems with the sbip framework. Technical Report TR-2017-6, Verimag Research Report, 2017.
11. G. Roşu and K. Havelund. Rewriting-based techniques for runtime verification. *Automated Software Engineering*, 12(2):151–197, Apr 2005.
12. H. L. S. Younes. *Verification and Planning for Stochastic Processes with Asynchronous Events*. PhD thesis, Carnegie Mellon, 2005.