Runtime Verification: The Application Perspective

Yliès Falcone, Lenore D. Zuck

► To cite this version:

Yliès Falcone, Lenore D. Zuck. Runtime Verification: The Application Perspective. International Journal on Software Tools for Technology Transfer, 2015, pp.121-123. 10.1007/978-3-642-34026-0_21 . hal-01127720

HAL Id: hal-01127720

<https://hal.science/hal-01127720>

Submitted on 7 Mar 2015

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.

Runtime Verification: the Application Perspective^{*}

Yliès Falcone¹, Lenore D. Zuck²

¹ University of Grenoble I (UJF), Laboratoire d'Informatique de Grenoble, e-mail: ylies.falcone@ujf-grenoble.fr

² University of Illinois at Chicago, e-mail: lenore@cs.uic.edu

The date of receipt and acceptance will be inserted by the editor

Abstract. In the past decade, Runtime Verification (RV) has gained much focus, from both the research community and practitioners. RV combines a set of theories, techniques and tools aiming towards efficient analysis of systems' executions and guaranteeing their correctness using monitoring techniques. Major challenges in RV include characterizing and formally expressing requirements that can be monitored, offering intuitive and concise specification formalisms, and monitoring specifications efficiently for functional and non-functional behavior. In spite of the major strides made in recent years, much effort is still needed to make RV an attractive and viable methodology for industrial use and to apply it to wider application domains, such as security, bio-health, power micro-grids.

This special issue of STTT proposes extended versions of four papers that have been selected from the Runtime Verification track at ISoLA 2012 [10].

during the system' execution. For these and many other reasons, RV offers an interesting alternative to static methods.

In the past decade, Runtime Verification (RV) has gained much focus from both research community and practitioners [13, 8, 12, 9, 4, 5]. RV combines a set of theories, techniques and tools aiming towards efficient analysis of systems' executions and guaranteeing their correctness using monitoring techniques. While some of the techniques used in RV are have been applied in several areas, mainly by the testing community, it had only recently become a first-class citizen in the formal-method community after a 2001 workshop (now a conference), carrying the name *runtime verification*, which was initiated by Klaus Havelund (who authors in this special issue of STTT) and Grigore Rosu.

Introduction

Static methods can guarantee program correctness. They are, however, not always applicable to a variety of systems and properties. Often the size of the system renders static methods prohibitively expensive. Systems to which static techniques are applied are those where correctness is to be proven under all circumstances, such as safety critical systems. In contrast, many "real-life" systems may be occasionally faulty, especially when the fault is not catastrophic (or even very expensive) and the system can recover from it. Similarly, static techniques are applicable to systems that are built top down and are often applied at the design stages. In contrast, many "real-life" systems are developed ad-hoc so that their properties are not always known à priori, yet, they may be learnt

This issue of STTT presents some new directions in RV. Two papers focus on the exploration of the expressiveness-efficiency spectrum [1], that is, the observed duality between the expressiveness of the specification language and the resource consumption made by runtime monitors. They support the common conjecture that the more expressive the specification formalism is, the more complex (and resource consuming) the associated monitoring algorithm is.

The other two papers focus on *statistical model-checking* (see [14]), which augments runtime verification with statistics. When one has access to several executions of the system under scrutiny, statistical model-checking allows to complement runtime verification of properties by using models of stochastic behavior to model the uncertainty of a system and to better assess the overall correctness of the system by applying statistical inference to reason on several execution traces.

^{*} The work of the second author was funded in part by NSF award CCF-0916438.

The Expressiveness/Efficiency Spectrum

On Piggyback Runtime Monitoring of Object-Oriented Programs (Hallé et al.)

The work in [6] proposes an original approach to the quest for expressive and efficient runtime monitors. It relies on the observation that when monitoring Java programs, many of the properties on the usage of data structures (usually addressed in benchmarks) are already monitored by the object instances at runtime. More precisely, the authors study the extent that the fields of an object carry information that can be piggybacked by a monitor to take a decision on the satisfaction of the property under verification. The authors also address the general question of how a monitor can use the information in the object's member fields. They empirically evaluate the benefit of piggyback monitoring and highlight the benefits of their approach.

Rule-based Runtime Verification Revisited (Havelund)

Inspired by the RETE algorithm ([2]), [7] proposes to revisit rule-based monitoring. While rule-based RV employs somewhat less efficient runtime monitors, it offers for a concise and elegant expression of specifications. In both RV and artificial intelligence, a rule is specified by a collection of facts, or events, that trigger some actions. The paper shows that the RETE algorithm can be made an efficient solution to the event-matching problem, and thus to the problem of (efficiently) dispatching events carrying data value to the related monitor instances. It also shows how to adapt and optimize RETE for monitoring purposes. An implementation of the modified RETE is proposed in the Scala-based tool LOGFIRE, whose performance is compared to state-of-the-art ruled-based systems.

Statistical Model Checking: Augmenting Runtime Verification with Statistics

Statistical Model Checking QoS Properties of Systems with SBIP (Nouri et al.)

The work in [11] proposes SBIP — a stochastic extension of the Behavior Interaction Priority (BIP) framework which is an expressive and rigorous component-based design flow for the hierarchical construction of systems. Adding stochastic features to BIP allows to better assess the correctness of the system as well as to model uncertainty stemming from faults or assumptions on the runtime platform. As [11] shows, SBIP allows to combine the results from different executions with statistical inference algorithms so as to add confidence measurement on the satisfaction of properties. To overcome the restrictions of using statistical model checking, the authors propose to consider properties expressed in Bounded Linear Temporal Logic and to eliminate the non-determinism

in the system by randomizing transitions. Two case studies conducted with SBIP are presented.

Scheduleability of Herschel-Planck Revisited Using Statistical Model Checking (David et al.)

The work in [3] proposes a scheduleability analysis of Herschel-Planck satellite system using symbolic model checking and statistical model checking. The paper shows the complementarity of these techniques for proving that a run is either realizable or that a deadline violation exists. It is demonstrated that statistical model checking improves performance analysis by providing response times when a system is scheduleable, and probability of deadline violation otherwise.

Acknowledgment

We would like to thank the organizing committees of ISoLA 2012 for setting up such a successful event, the programme committee and reviewers of the conference and the special issue of STTT for helping with the selection of papers, and all authors who contributed to the track. We would especially like to thank the authors for providing us with such excellent papers, and the referees for their diligent work.

References

1. Howard Barringer, Yliès Falcone, Klaus Havelund, Giles Reger, and David Rydeheard. Quantified event automata: Towards expressive and efficient runtime monitors. In *FM 2012: 18th International symposium on Formal Methods, Paris, France, August 27-31, 2012*, volume 7436 of *Lecture Notes in Computer Science*, pages 65–79, 2012.
2. Howard Barringer, David E. Rydeheard, and Klaus Havelund. Rule systems for run-time monitoring: from Eagle to RuleR. *J. Log. Comput.*, 20(3):675–706, 2010.
3. Alexandre David, Kim G. Larsen, Axel Legay, and Marius Mikučionis. Scheduleability of herchel-planck revisited using statistical model checking. In *STTT*, 2014.
4. Yliès Falcone, Jean-Claude Fernandez, and Laurent Mounier. Runtime verification of safety-progress properties. In Saddek Bensalem and Doron Peled, editors, *Runtime Verification, 9th International Workshop, RV 2009, Grenoble, France, June 26-28, 2009. Selected Papers*, volume 5779 of *Lecture Notes in Computer Science*, pages 40–59. Springer, 2009.
5. Yliès Falcone, Klaus Havelund, and Giles Reger. A tutorial on runtime verification. In Manfred Broy, Doron Peled, and Georg Kalus, editors, *Engineering Dependable Software Systems*, volume 34 of *NATO Science for Peace and Security Series, D: Information and Communication Security*, pages 141–175. IOS Press, 2013.
6. Sylvain Hallé, Jason Vallet, and Raphaël Tremblay-Lessard. On piggyback runtime monitoring of object-oriented programs. In *STTT*, 2014.
7. Klaus Havelund. Rule-based runtime verification revisited. In *STTT*, 2014.

8. Klaus Havelund and Allen Goldberg. Verify your runs. In Bertrand Meyer and Jim Woodcock, editors, *VSTTE*, volume 4171 of *Lecture Notes in Computer Science*, pages 374–383. Springer, 2005.
9. Martin Leucker and Christian Schallhart. A brief account of runtime verification. *Journal of Logic and Algebraic Programming*, 78(5):293–303, may/june 2008.
10. Tiziana Margaria and Bernhard Steffen, editors. *Proceedings of the 5th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation, ISOLA 2012, Amirandes, Heraclion, Crete, October 15-18, 2012*, Lecture Notes in Computer Science. Springer, 2012.
11. Ayoub Nouri, Saddek Bensalem, Benoit Delahaye, Marius Bozga, Cyrille Jegourel, and Axel Legay. Statistical model checking qos properties of systems with SBIP. In *STTT*, 2014.
12. Amir Pnueli and Aleksandr Zaks. PSL Model Checking and Run-Time Verification Via Testers. In Jayadev Misra, Tobias Nipkow, and Emil Sekerinski, editors, *FM*, volume 4085 of *Lecture Notes in Computer Science*, pages 573–586. Springer, 2006.
13. Runtime Verification. <http://www.runtime-verification.org>, 2001-2012.
14. Høakan L. S. Younes. *Verification and Planning for Stochastic Processes with Asynchronous Events*. PhD thesis, Carnegie Mellon, 2005.