Runtime Verification - 17 Years Later

Klaus Havelund[1(✉)] and Grigore Roşu[2]

[1] Jet Propulsion Laboratory, California Institute of Technology, Pasadena, USA
klaus.havelund@jpl.nasa.gov
[2] University of Illinois at Urbana-Champaign, Urbana, USA

Abstract. Runtime verification is the discipline of analyzing program executions using rigorous methods. The discipline covers such topics as specification-based monitoring, where single executions are checked against formal specifications; predictive runtime analysis, where properties about a system are predicted/inferred from single (good) executions; specification mining from execution traces; visualization of execution traces; and to be fully general: computation of any interesting information from execution traces. Finally, runtime verification also includes fault protection, where monitors actively protect a running system against errors. The paper is written as a response to the 'Test of Time Award' attributed to the authors for their 2001 paper [45]. The present paper provides a brief overview of what lead to that paper, what has happened since, and some perspectives on the future of the field.

1 Introduction

Runtime verification (RV) [10,26,39,55] has emerged as a field of computer science within the last couple of decades. RV is concerned with the rigorous monitoring and analysis of software and hardware system executions. The field, or parts of it, can be encountered under several other names, including, e.g., runtime checking, monitoring, dynamic analysis, and runtime analysis. Since only single executions are analyzed, RV scales well compared to more comprehensive formal methods, but of course at the cost of coverage. Nonetheless, RV can be useful due to the rigorous methods involved. Conferences and workshops are now focusing specifically on this subject, including the Runtime Verification conference, which was initiated by the authors in 2001 as a workshop and became a conference in 2010, and runtime verification is now also often listed as a subject of interest in other conference calls for papers.

The paper is written as a response to the 'Test of Time Award' attributed to the authors for the 2001 paper [45] (*Monitoring Java Programs with Java PathExplorer*), presented 17 years ago (at the time of writing) at the first Runtime Verification workshop (RV'01) in Paris, July 23, 2001.

K. Havelund—The research performed by this author was carried out at Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

C. Colombo and M. Leucker (Eds.): RV 2018, LNCS 11237, pp. 3–17, 2018.
https://doi.org/10.1007/978-3-030-03769-7_1

This paper reports on our own RV work, with some references to related work that specifically inspired us, and discusses the lessons learned and our perspective on the future of this field. Note that we do not try to identify all literature that inspired us. That task would be impossible. Previous publications of ours [26,42,44] have provided more technical tutorial-like presentations of the field. This paper rather offers information about the motivations for our work and philosophical considerations. As such this paper is closer in spirit to the longer paper [43]. It should be mentioned that most of the works over time have been done in collaboration with other people and inspired/initiated/driven by other people. We have just been lucky to be in the midst of all this work.

The paper is organized according to the time line of events, first leading up to [45], then the work described in that paper, the work that followed, and finally some thoughts on the future of this field.

2 In the Beginning

The initial interest of the first author in formal methods stems from his involvement in the design of the RAISE specification language RSL [30], during the period 1984–1991, and even with earlier work in the early 1980's on developing a parser and type checker for its predecessor VDM [14,15,28]. These are so-called wide-spectrum specification languages permitting formal specification at a high level, and "programming" at a low level, all within the same language, supported by a formal refinement relation between the different levels. These languages were impressively ahead of their time if one looks at these from a programming language perspective. For example, VDM^{++} has many similarities with today's SCALA programming language.

However, these languages were fundamentally still specification languages, and not programming languages, in spite of the fact that these languages have a lot in common with modern high-level programming languages, such as e.g. ML. The thought therefore was: why not benefit from the evolution of modern high-level programming languages and focus on verification of such? This was the first step: the focus on programs rather than models. This lead to the work [34] of the first author on attempting to develop a specification language for an actual programming language, namely CONCURRENT ML (CML), an extension of Milner's ML with concurrency.

Later work with the very impressive PVS theorem prover [35] helped realize that theorem proving is hard after all, and that some form of more automated reasoning on programs would be useful as a less perfect alternative. Hence, thus far the realization was that *automated* verification of *programs* was a desirable objective. Note that at the time the main focus in the formal methods community was on models, not programs.

The next big move was the development of the JAVA PATHFINDER (JPF), a JAVA model checker, first as a translator from JAVA to the PROMELA modeling language of the SPIN model checker [41] (often referred to as JPF1), and later as a byte code model checker [50] (occasionally referred to as JPF2). The goal of

this work was to explore how far model checking could be taken wrt. real code verification, either using JAVA as just a better modeling language, or, in the extreme case, for model checking real programs. A sub-objective was to explore the space between testing and full model checking.

JPF1 suffered from the problem of translating a complex language such as JAVA to the much simpler language PROMELA, resulting in a sensation that this approach worked for some programs but not for all programs. It was hard to go the last 20%. JPF2 solved part of this problem, but still suffered from the obvious problem of state space explosion. In addition, the model checker itself was a homemade JVM on top of the real JVM, and hence slow.

At this time we came across two inspiring invited talks at the SPIN 2000 workshop, which we organized. The first was a presentation by Harrow from Compaq on the VISUALTHREADS tool [33]. The purpose of this tool was to support Compaq's customers in avoiding multithreading errors. Specifically two algorithms appeared interesting: predictive data race and deadlock detection. These algorithms can detect the *potential* for a data race or deadlock by analyzing a run that does not necessarily encounter the error. The second invited talk was presented by Drusinsky, on the TEMPORAL ROVER [25] for monitoring temporal logic properties. We implemented the data race algorithm, also known as the Eraser algorithm [61], and a modification of the deadlock detection algorithm in JPF2. The idea was to first execute the program to check for data races and deadlocks using the two very scalable algorithms, and then only if error potentials were found between identified threads, to launch the model checker focusing specifically on those threads.

The two authors of [45] met at NASA Ames in 2000, when the second author started his first job right out of school, and this way, without knowing it, a fruitful, life-time collaboration and friendship with the first author. Inspired by recent joint work with his PhD adviser, Joseph Goguen, the second author was readily convinced that otherwise heavy-weight specification-based analysis techniques can very well apply to execution traces instead of whole systems, and thus achieve scalability by analyzing only what happens at runtime, as it happens. This, paired with provably correct recovery, gives the same level of assurance as formal verification of the whole system, but in a manner that allows us to divide-and-conquer the task. So the second author was "all in", ready to use his fresh algebraic specification and formal verification knowledge to rigorously analyze execution traces.

At this point, the previously mentioned observations about scalability of the traditional verification approaches, the experiments with data race and deadlock detection mentioned above, and some other less technical issues, led to our research focusing just on observing program executions. A constraint was that it should not be based on test case generation, since so many people were studying this already. We wanted to follow the path less explored. This is where the JAVA PATHEXPLORER project began, inspired by other work, but not too much other work.

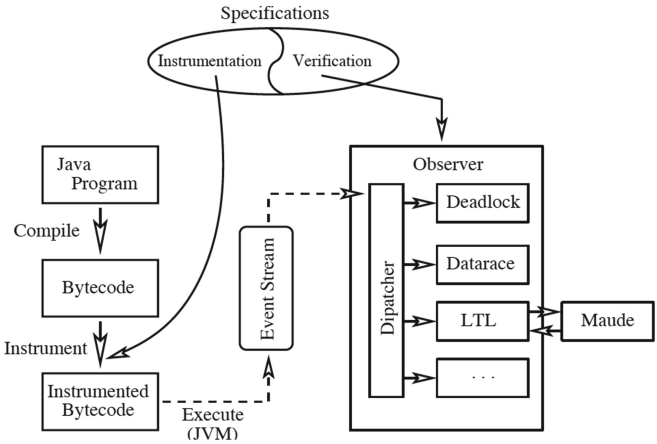

Fig. 1. The JPaX architecture.

3 Java PathExplorer

Our first pure runtime verification system was Java PathExplorer (JPaX), described in the award winning paper [45], as well as in other papers [46–49,60]. The system is briefly described below.

3.1 Architecture

JPaX was a general framework for analyzing execution traces. It supported three kinds of algorithms: propositional temporal logic conformance checking, data race detection, and deadlock detection, as discussed earlier. Figure 1 shows JPaX's architecture. A Java program is instrumented (at byte code level) to issue events to the monitoring side, which is customizable, allowing the addition of new monitors. The temporal logic monitoring module was originally based on a propositional future time linear temporal logic, but was later extended to also cover past time.

An interesting aspect of the system was the use of the Maude [21] rewriting system for implementing monitoring logics as deep DSLs. One could in very few lines implement, e.g., linear temporal logic (LTL), with syntax and its monitoring algorithm, and have Maude function as the monitoring engine as well. There was a grander vision present at the time: to use a powerful Turing complete language, such as Maude, for monitoring, and not be restricted to just, e.g., LTL. However, that vision did not evolve beyond the thought stage, and had to wait some additional years, as discussed in Sect. 4. Below we briefly discuss some of the algorithms developed during the JPaX project.

Future Time LTL. The future time LTL monitoring used Maude to rewrite formulas. Consider, e.g., the LTL formula $p\,\mathcal{U}\,q$, meaning q eventually becomes

true and until then p is true. The implementation of JPaX was based on classical equational laws for temporal operators, such as:

$$p \,\mathcal{U}\, q = q \wedge \bigcirc(p \,\mathcal{U}\, q) \qquad \text{and} \qquad \Box p = p \wedge \bigcirc(\Box p) \tag{1}$$

Consider the sample formula $\Box(green \rightarrow \bigcirc(\neg red \,\mathcal{U}\, yellow))$. Upon encountering a *green* in a trace, the formula will be rewritten into the following formula, which must be true in the next state: $(\neg red \,\mathcal{U}\, yellow) \wedge \Box(green \rightarrow (\neg red \,\mathcal{U}\, yellow))$. In Maude this was realized by a few simple rewrite rules, including the following two for the until operator (E is an event and T is a trace, the first rule handles the case of a trace consisting of only one event):

```
eq E |= X U Y = E |= Y.
eq E,T |= X U Y = E,T |= Y or E,T |= X and T |= X U Y.
```

3.2 Past Time LTL

Later, an efficient dynamic programming algorithm for monitoring *past time* linear temporal logic was developed [48], inspired by an initial encoding in Maude described in [45]. Consider the following past time formula: $red \rightarrow \blacklozenge green$ (whenever *red* is observed, in the past there has been a *green*). The algorithm for checking past time formulas like this uses two arrays, now and pre, recording the status of each sub-formula now and previously. Index 0 refers to the formula itself with positions ordered by the sub-formula relation. Then for this property, for each observed event the arrays are updated as follows.

```
bool pre [0..3], now [0..3];

fun processEvent(e) {                  // Sub−formula:
  now[3] := (event = red)              // red
  now[2] := (event = green)            // green
  now[1] := now[2] || pre[1]           // PREV green
  now[0] := !now[3] || now[1]          // red −> PREV green
  if !now[0] then output(" property violated ");
  pre := now;
}
```

This dynamic programming algorithm was generalized and optimized in [49,59] and later found way into three other systems for monitoring parametric temporal formulas, namely MOP [57], MonPoly [11], and DejaVu [40].

3.3 Data Races and Deadlocks

When used for bug finding, the effectiveness of runtime verification depends on the choice of test suite. For concurrent systems this is critical, due to the many possible non-deterministic execution paths. *Predictive runtime verification* approaches this problem by replacing a target property P with a stronger

property Q such that there is a high probability that the program satisfies P iff a random trace of the program will satisfy Q. As already mentioned, one such algorithm was the Eraser algorithm [61], for detecting *potentials* for data races (where two threads can access a shared variable simultaneously). It is often referred to as the *lock set* algorithm as each variable is associated with a set of locks protecting it. The *lock graph* algorithm [33], would detect "dining philosopher"-like deadlock potentials by building a simple lock graph where a cycle indicates a deadlock potential. In [13] we augmented the original lock graph algorithm to reduce false positives in the presence of so-called guard locks (locks that prevent cyclic deadlocks). That paper was later followed by [12], which suggested a code instrumentation method (inserting wait statements) for confirming found deadlock potentials. Other forms of data races than those detected by Eraser are possible. In [3] a dynamic algorithm for detecting so-called high-level data races (races involving collections of variables) is described.

3.4 Code Instrumentation

JPaX code instrumentation was performed with Compaq's JTrek [22], a Java byte code instrumentation tool. Operating at the byte code level offers expressive power, but makes writing code instrumentation instructions inconvenient. An attempt was later made to develop an easier to use code instrumentation tool named JSpy [31] on top of JTrek. In this tool code instrumentation could be expressed as a set of high-level rules, formulated in Java (an internal Java DSL), each consisting of a predicate and an action.

3.5 Trace Visualization

Execution trace visualization is a subject that in our opinion has promising potential, although our own involvement in this direction is limited to [4]. The advantage of visualization is that it can provide a free-of-charge abstract view of the trace, from which a user potentially may be able to conclude properties about the program, or at least the execution, without having to explicitly formulate these properties. We can distinguish between two forms of trace visualization as outlined in [4]: *still visualization*, where all events are visualized in one view, and *animated visualization*. In [4], an extension of Uml sequence diagrams with symbols is described for representing still visualizations of the execution of concurrent programs.

4 The Aftermath

The period after JPaX followed two tracks, which can be summarized as: experiments with aspect-oriented programming for program instrumentation, and so-called parametric monitoring of events carrying data.

4.1 Aspect-Oriented Programming

Whilst initial runtime verification frameworks targeted Java, the RMOR (Requirement Monitoring and Recovery) framework [36] targeted the monitoring of C programs against state machines using a homegrown aspect-oriented framework to perform program instrumentation. RMOR was implemented in OCAML using CIL (C Intermediate Language), a C program analysis and transformation system, itself written in OCAML. Later it was attempted to "go all aspect-oriented", meaning that aspects no longer were thought of as just the plumbing for performing code instrumentation, but instead that monitors *are* aspects. Some of our experiments went in the direction of what today is called *state-full aspects* [1,65]. Here one takes a starting point in an aspect-oriented language framework (such as e.g. ASPECTJ) and extends it with so-called *tracecuts*, denoting predicates on the execution trace. An advice can be associated with a tracecut, and executes when the tracecut is matched by the execution. We proposed this line of work already in [27]. Other later work included [16,51,62,63]. The main observation in these works was that aspect-oriented programming can be extended vertically (allowing more pointcuts) and horizontally (allowing temporal advice, essentially monitoring temporal constraints).

4.2 Runtime Verification with Data

JPAX had a number of limitations. The perhaps most important was the propositional nature of the temporal logics. One could not, for example, monitor parametric events carrying data, such as *openFile("data.txt")*, where *openFile* is an event name and *"data.txt"* is data. It is perhaps of interest to note, that at the time we were not (and are still not) aware of any system that at the time was able to monitor such parametric events in a temporal logic.

4.3 The Beginning of Data

These considerations lead to two different systems: EAGLE [6] and MOP [19]. EAGLE was a small and general logic having similarities with a linear time μ-calculus, supporting monitoring events with data, and allowing user-defined temporal operators. The later HAWK system [23] was an attempt to tie EAGLE to the monitoring of JAVA programs with automated code instrumentation using aspect-oriented programming, specifically ASPECTJ [53].

The same JPAX limitations that motivated the development of EAGLE also stimulated the apparition of monitoring-oriented programming (MOP) [18–20]. MOP proposed that runtime monitoring be supported and encouraged as a fundamental principle of software development, where monitors are automatically synthesized from formal specifications and integrated at appropriate places in the program. Violations and/or validations of specifications can trigger user-defined code at any points in the program, in particular recovery code, outputting/sending messages, or raising exceptions. MOP has made three important early contributions. First, it proposed specification formalism independence, allowing users

to insert their favorite or domain-specific requirements specification formalisms via *logic plugin* modules. Second, it proposed automated code instrumentation as a means to weave the monitoring checking code within the application; the first version in 2003 used Perl for instrumentation [19], while the subsequent versions starting with 2004 [18] used AspectJ [53]. Finally, it proposed a formalism-independent semantics and implementation for parametric specifications. Conceptually, execution traces are sliced according to each observed instance of the parameters, and each slice is checked by its own monitor instance in a manner that is independent of the employed specification formalism. The practical challenge is how to deal with the potentially huge number of monitor instances. JavaMop proposed several optimizations, presented in [58] together with the mathematical foundations of parametric monitoring.

The Eagle system mentioned earlier was considered quite an elegant system, but its implementation was complicated. The subsequent rule-based lower level Ruler system [9] was meant as an "assembler" into which other temporal specification languages could be compiled for efficient trace checking. However, it assumed a life of its own as a specification language. Ruler was given a finite-trace semantics with four verdicts. The verdicts STILL_TRUE and STILL_FALSE are given if the rule system would accept/reject the trace if it were to end at the current event, whilst the verdicts TRUE and FALSE were reserved for traces where every extension would be accepted/rejected. Ruler allowed for very complex rule systems that could be *chained* together such that one rule system produced outputs for another rule system to consume as input events. Rule systems could be combined sequentially, in parallel, and conditionally.

A project solidly rooted in an actual space mission was the development of the LogScope temporal logic for log analysis [7]. The purpose of the project was to assist the team testing the flight software for JPL's Mars rover Curiosity, which successfully landed on Mars on August 6, 2012. The software produces rich log information. Traditionally, these logs are analyzed with complex Python scripts. The LogScope logic was developed to support notations comprehensible to test engineers, including a very simple and convenient data parameterized temporal logic, which was translated to a form of data parameterized automata, which themselves could be used for specification of more complex properties that the temporal logic could not express. LogScope was furthermore implemented in Python, allowing Python code fragments to be included in specifications, all in order to integrate with the existing Python scripting culture at JPL.

4.4 Internal DSLs

Earlier we mentioned a grander vision to use a powerful Turing complete language for monitoring. The fundamental problem with a logic is that it likely may be insufficient for practical purposes if not designed extremely optimally. Engineers are, e.g., often observed using Python for monitoring tasks. Of course in lack of a better notation, but also because it provides expressive power to perform arbitrary computations, e.g. on observed data. This observation led to several

experiments with so-called internal DSLs, where one extends a programming language with monitoring features. This allows the user to use the features of the programming language when the features of the monitoring logic do not suffice. TraceContract [8,37] is such an internal Scala DSL (effectively an API) for monitoring, based on a mixture of temporal logic and state machines. It is developed using Scala's features for defining internal DSLs. TraceContract, although a research tool, was later used for analysis of command sequences sent to NASA's LADEE (Lunar Atmosphere and Dust Environment Explorer) spacecraft throughout its mission.

Another example of an internal Scala DSL is LogFire [38]. LogFire is a rule-based system similar to Ruler, but based on a modification of the Rete algorithm [24,29], used in several rule-based systems. LogFire was part of an investigation of the Rete algorithm's applicability for runtime verification. LogFire has become part of the software that daily processes telemetry data from JPL's Mars Curiosity rover. LogFire's ability to generate facts can be used for Complex Event Processing (CEP) [56], where higher-level events (abstractions) are generated from lower-level events. CEP can be used for further analysis and/or human comprehension, e.g. through visualization. Another CEP system is Nfer [52], which in part was influenced by our work on rule-based systems, and LogFire in particular. The result of applying an Nfer specification to an event stream is a set of time bounded intervals. The specification consists of rules of the form: name :− body (a rule name followed by a rule body). The semantics is similar to that of Prolog (hence the :− symbol): when the body is true an interval is generated with that name. A difference from Prolog is that rule bodies contain temporal constraints based on operators from Allen Temporal Logic [2]. Nfer was created due to a need for comprehending large telemetry streams from Mars rovers. Abstracting these to higher level intervals, compared to the low level raw event stream, should ease human comprehension.

4.5 First-Order Beyond Slicing

Ruler, as a layer of syntactic sugar on top of the rule formalism, offered a sub-formalism resembling a data parameterized automaton language. Likewise, LogScope, inspired by Ruler, offered a data parameterized automaton notation (in addition to the temporal logic). Quantified event automata (Qea) [5] was an attempt to design a pure data parameterized automaton monitoring system logic, using the efficient trace slicing approach previously introduced in the JavaMop tool [57], but dealing with some of the limitations with respect to expressiveness. A Qea specification consists of a list of first-order quantifications (universal and existential) and an automaton. They can be compared to extended state machines (allowing arbitrary guards and actions on transitions operating on local state, but are more succinct due to the fact that automata are "spawned" according to parameters (there is a local state for each combination of parameters).

A different approach to optimizing monitoring of parametric data is implemented in the DejaVu tool [40], which uses BDDs [17] to efficiently represent

data observed in the trace. Logic-wise, the system supports a standard past time temporal logic with quantification. The logic in itself is not the innovation, rather it is the use of BDDs to represent the sets of values observed in the trace for the quantified variables. The representation of sets of assignments as BDDs allows a very simple algorithm that naturally extends the dynamic programming monitoring algorithm for propositional past time temporal logic shown on page 5 and presented in [47], using two vectors *now* and *pre*. However, while in [47] the vectors contain Boolean values, here the values are BDDs.

5 Discussion

Numerous runtime verification logics have been developed over time. They include various forms of temporal logics, state machines, regular expressions, context free grammars, rule systems, variations of the μ-calculus, process algebras, stream processing, timed versions of these, and even statistical versions, where data can be computed as part of monitoring. It is clear that parametric/first-order versions of these logics are needed. Some efforts have been made to combine two or more of these logics, such as, e.g., combining temporal logic and regular expressions. An interesting trend is logics which not just produce a Boolean value, but rather a data value of any type. This leads to systems computing arbitrary data values from traces. It is, however, nearly impossible at this point to estimate which of these approaches would potentially get infused in industrial settings.

Whether to develop a DSL as external or internal is a non-trivial decision. An external DSL is usually cleaner and more directly tuned towards the immediate needs of the user. In addition, they are easier to process and therefore optimize for efficiency. However, the richer the DSL becomes (moving towards Turing-completeness) the harder the implementation effort becomes. An internal DSL can be very fast to implement and augment with new (even user-defined) operators, and can provide an expressiveness that would require a major effort to support in an external DSL. One also gains the advantage of IDEs for the host language. A hypothesis is that monitoring logics used in practice will need to support very expressive expression languages to process data, such as strings and numbers that are part of the observed events. Temporal logic could become part of a programming language assertion language. This could be seen as part of a design-by contract approach also supporting pre/post conditions and class invariants.

An important topic may be inferring specifications from execution traces. Our own limited work in this area includes [54,64]. Related to specification mining is execution trace visualization (the visualization can be considered a learned model). The advantage of visualization is that it can provide a free-of-charge abstract view of the trace, from which a user potentially may be able to conclude properties about the program, or at least the execution, without having to explicitly formulate these properties.

Full verification is of course preferred over partial verification performed by a monitor. The combination of static and dynamic verification can provide the best

of both worlds: prove as much as is feasible statically and verify the remaining proof obligations during runtime. To properly achieve this goal, we need formal specifications not only for the properties to verify, but also for the programming language itself. Moreover, we need provably correct monitor generation techniques, so we can put all the specification and proof artifacts together and assemble a proof of correctness for the entire system. Interestingly, once a specification of the programming language itself is available, then one can go even one step further and monitor the execution of the program even against the language specification. This may seem redundant at first, but it actually makes full sense for some languages with complex semantics, like C. For example, tools like Valgrind or UBSan detect undefined behaviors in C/C++ programs, which are essentially deviations from the intended language semantics. The RV-Match tool [32] is an attempt to push runtime verification in this direction.

In fault-protection strategies, the goal is to recover the system once it has failed. The general problem of how to recover from a bad program state is interesting and quite challenging. The ultimate solution to this problem can be found in planning and scheduling systems, where a planner creates a plan (straight-line program) to execute for a limited time period, an executive executes the plan, and a monitor monitors the execution. Upon failure detected by the monitor, a new plan (program) is generated online.

References

1. Allan, C., et al.: Adding trace matching with free variables to AspectJ. SIGPLAN Not. **40**, 345–364 (2005)
2. Allen, J.F.: Maintaining knowledge about temporal intervals. Commun. ACM **26**(11), 832–843 (1983)
3. Artho, C., Havelund, K., Biere, A.: High-level data races. Softw. Test. Verif. Reliab. **13**(4), 207–227 (2004)
4. Artho, C., Havelund, K., Honiden, S.: Visualization of concurrent program executions. In: 31st Annual International Computer Software and Applications Conference (COMPSAC 2007), vol. 2, pp. 541–546, July 2007
5. Barringer, H., Falcone, Y., Havelund, K., Reger, G., Rydeheard, D.: Quantified event automata: towards expressive and efficient runtime monitors. In: Giannakopoulou, D., Méry, D. (eds.) FM 2012. LNCS, vol. 7436, pp. 68–84. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32759-9_9
6. Barringer, H., Goldberg, A., Havelund, K., Sen, K.: Rule-based runtime verification. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 44–57. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24622-0_5
7. Barringer, H., Groce, A., Havelund, K., Smith, M.: Formal analysis of log files. J. Aerosp. Comput. Inf. Commun. **7**(11), 365–390 (2010)
8. Barringer, H., Havelund, K.: TraceContract: a Scala DSL for trace analysis. In: Butler, M., Schulte, W. (eds.) FM 2011. LNCS, vol. 6664, pp. 57–72. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21437-0_7
9. Barringer, H., Rydeheard, D.E., Havelund, K.: Rule systems for run-time monitoring: from Eagle to RuleR. J. Log. Comput. **20**(3), 675–706 (2010)

10. Bartocci, E., Falcone, Y., Francalanza, A., Reger, G.: Introduction to runtime verification. In: Bartocci, E., Falcone, Y. (eds.) Lectures on Runtime Verification. LNCS, vol. 10457, pp. 1–33. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75632-5_1
11. Basin, D., Klaedtke, F., Müller, S., Pfitzmann, B.: Runtime monitoring of metric first-order temporal properties. In: Proceedings of the 28th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, volume 2 of Leibniz International Proceedings in Informatics (LIPIcs), pp. 49–60. Schloss Dagstuhl - Leibniz Center for Informatics (2008)
12. Bensalem, S., Fernandez, J.-C., Havelund, K., Mounier, L.: Confirmation of deadlock potentials detected by runtime analysis. In: Parallel and Distributed Systems: Testing and Debugging (PADTAD 2006), Portland, Maine, USA, July 2006
13. Bensalem, S., Havelund, K.: Dynamic deadlock analysis of multi-threaded programs. In: Ur, S., Bin, E., Wolfsthal, Y. (eds.) HVC 2005. LNCS, vol. 3875, pp. 208–223. Springer, Heidelberg (2006). https://doi.org/10.1007/11678779_15
14. Bjørner, D., Jones, C.B. (eds.): The Vienna Development Method: The Meta-Language. LNCS, vol. 61. Springer, Heidelberg (1978). https://doi.org/10.1007/3-540-08766-4
15. Bjørner, D., Jones, C.B.: Formal Specification and Software Development. Prentice Hall International (1982). ISBN 0-13-880733-7
16. Bodden, E., Havelund, K.: Aspect-oriented race detection in Java. IEEE Trans. Softw. Eng. **36**(4), 509–527 (2010)
17. Bryant, R.E.: Symbolic Boolean manipulation with ordered binary-decision diagrams. ACM Comput. Surv. (CSUR) **24**(3), 293–318 (1992)
18. Chen, F., D'Amorim, M., Roşu, G.: A formal monitoring-based framework for software development and analysis. In: Davies, J., Schulte, W., Barnett, M. (eds.) ICFEM 2004. LNCS, vol. 3308, pp. 357–372. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-30482-1_31
19. Chen, F., Roşu, G.: Towards monitoring-oriented programming: A paradigm combining specification and implementation. In: Proceedings of the 3rd International Workshop on Runtime Verification (RV 2003), volume 89(2) of Electronic Notes Theoretical Computer Science, pp. 108–127. Elsevier Science Inc. (2003)
20. Chen, F., Roşu, G.: MOP: an efficient and generic runtime verification framework. In: Object-Oriented Programming, Systems, Languages and Applications (OOPSLA 2007), pp. 569–588. ACM, ACM SIGPLAN Notices (2007)
21. Clavel, M., et al.: Maude: specification and programming in rewriting logic. Theor. Comput. Sci. **285**(2), 187–243 (2002)
22. Cohen, S.: JTrek. (2001)
23. d'Amorim, M., Havelund, K.: Event-based runtime verification of Java programs. ACM SIGSOFT Softw. Eng. Notes **30**(4), 1–7 (2005)
24. Doorenbos, R.B.: Production Matching for Large Learning Systems. Ph. D. thesis, Carnegie Mellon University, Pittsburgh, PA (1995)
25. Drusinsky, D.: The temporal rover and the ATG rover. In: Havelund, K., Penix, J., Visser, W. (eds.) SPIN 2000. LNCS, vol. 1885, pp. 323–330. Springer, Heidelberg (2000). https://doi.org/10.1007/10722468_19
26. Falcone, Y., Havelund, K., Reger, G.: A tutorial on runtime verification. In: Broy, M., Peled, D., Kalus, G., (eds.) Engineering Dependable Software Systems, volume 34 of NATO Science for Peace and Security Series - D: Information and Communication Security, pp. 141–175. IOS Press (2013)

27. Filman, R., Havelund, K.: Source-code instrumentation and quantification of events. In: Foundations of Aspect-Oriented Languages (FOAL 2002), Enschede, The Netherlands, April 2002
28. Fitzgerald, J., Larsen, P.G., Mukherjee, P., Plat, N., Verhoef, M.: Validated Designs for Object-oriented Systems. Springer, TELOS, Santa Clara (2005)
29. Forgy, C.: Rete: a fast algorithm for the many pattern/many object pattern match problem. Artif. Intell. **19**, 17–37 (1982)
30. George, C., et al.: The RAISE Specification Language. The BCS Practitioner Series. Prentice-Hall, Hemel Hampstead (1992)
31. Goldberg, A., Havelund, K.: Instrumentation of Java bytecode for runtime analysis. In: Fifth ECOOP Workshop on Formal Techniques for Java-like Programs (FTfJP 2003), Darmstadt, Germany, July 2003
32. Guth, D., Hathhorn, C., Saxena, M., Roşu, G.: RV-Match: practical semantics-based program analysis. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016, Part I. LNCS, vol. 9779, pp. 447–453. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-41528-4_24
33. Harrow, J.J.: Runtime checking of multithreaded applications with visual threads. In: Havelund, K., Penix, J., Visser, W. (eds.) SPIN 2000. LNCS, vol. 1885, pp. 331–342. Springer, Heidelberg (2000). https://doi.org/10.1007/10722468_20
34. Havelund, K.: The Fork Calculus - Towards a Logic for Concurrent ML. Ph. D. thesis. DIKU, Department of Computer Science, University of Copenhagen, Denmark (1994)
35. Havelund, K.: Mechanical verification of a garbage collector. In: Rolim, J., et al. (eds.) IPPS 1999. LNCS, vol. 1586, pp. 1258–1283. Springer, Heidelberg (1999). https://doi.org/10.1007/BFb0098007
36. Havelund, K.: Runtime verification of C programs. In: Suzuki, K., Higashino, T., Ulrich, A., Hasegawa, T. (eds.) FATES/TestCom -2008. LNCS, vol. 5047, pp. 7–22. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-68524-1_3
37. Havelund, K.: Data automata in Scala. In: Proceedings of the 8th International Symposium on Theoretical Aspects of Software Engineering (TASE 2014). IEEE Computer Society (2014)
38. Havelund, K.: Rule-based runtime verification revisited. Int. J. Softw. Tools Technol. Trans. **17**(2), 143–170 (2015)
39. Havelund, K., Goldberg, A.: Verify your runs. In: Meyer, B., Woodcock, J. (eds.) VSTTE 2005. LNCS, vol. 4171, pp. 374–383. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69149-5_40
40. Havelund, K., Peled, D.A., Ulus, D.: First order temporal logic monitoring with BDDs. In: Formal Methods in Computer Aided Design (FMCAD), pp. 116–123. IEEE (2017)
41. Havelund, K., Pressburger, T.: Model checking Java programs using Java PathFinder. Int. J. Softw. Tools Technol. Transf. **2**(4), 366–381 (2000)
42. Havelund, K., Reger, G.: Runtime verification logics - a language design perspective. In: Aceto, L., Bacci, G., Bacci, G., Ingólfsdóttir, A., Legay, A., Mardare, R. (eds.) Models, Algorithms, Logics and Tools. LNCS, vol. 10460, pp. 310–338. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-63121-9_16
43. Havelund, K., Reger, G., Roşu, G.: Runtime verification - past experiences and future projections. volume 10000 of LNCS. Springer (2018)
44. Havelund, K., Reger, G., Thoma, D., Zălinescu, E.: Monitoring events that carry data. In: Bartocci, E., Falcone, Y. (eds.) Lectures on Runtime Verification. LNCS, vol. 10457, pp. 61–102. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-75632-5_3

45. Havelund, K., Roşu, G.: Monitoring Java programs with Java PathExplorer. In: Proceedings of the 1st International Workshop on Runtime Verification (RV 2001), vol. 55(2) of Electronic Notes Theoretical Computer Science. Elsevier, Paris, France, 23 July 2001. Won the RV 2018 Test of Time Award
46. Havelund, K., Roşu, G.: Monitoring programs using rewriting. In: Proceedings of the 16th IEEE International Conference on Automated Software Engineering (ASE 2001), pp. 135–143 (2001)
47. Havelund, K., Roşu, G.: Synthesizing monitors for safety properties. In: Katoen, J.-P., Stevens, P. (eds.) TACAS 2002. LNCS, vol. 2280, pp. 342–356. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-46002-0_24
48. Havelund, K., Roşu, G.: An overview of the runtime verification tool Java PathExplorer. Form. Methods Syst. Des. **24**(2), 189–215 (2004)
49. Havelund, K., Roşu, G.: Efficient monitoring of safety properties. Int. J. Softw. Tools Technol. Transf. **6**(2), 158–173 (2004)
50. Havelund, K., Visser, W.: Program model checking as a new trend. STTT **4**(1), 8–20 (2002)
51. Havelund, K., Wyk, E.V.: Aspect-oriented monitoring of C programs. In: The Sixth IARP-IEEE/RAS-EURON Joint Workshop on Technical Challenges for Dependable Robots in Human Environments, Pasadena, CA, May 17–18 2008
52. Kauffman, S., Havelund, K., Joshi, R.: nfer – a notation and system for inferring event stream abstractions. In: Falcone, Y., Sánchez, C. (eds.) RV 2016. LNCS, vol. 10012, pp. 235–250. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46982-9_15
53. Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., Griswold, W.G.: An overview of AspectJ. In: Knudsen, J.L. (ed.) ECOOP 2001. LNCS, vol. 2072, pp. 327–354. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-45337-7_18
54. Lee, C., Chen, F., Roşu, G.: Mining parametric specifications. In: Proceedings of the 33rd International Conference on Software Engineering, ICSE 2011, Waikiki, Honolulu, HI, USA, May 21–28 2011, pp. 591–600 (2011)
55. Leucker, M., Schallhart, C.: A brief account of runtime verification. J. Log. Algebr. Program. **78**(5), 293–303 (2008)
56. Luckham, D. (ed.): The Power of Events: An Introduction to Complex Event Processing in Distributed Enterprise Systems. Addison-Wesley, Boston (2002)
57. Meredith, P., Jin, D., Griffith, D., Chen, F., Roşu, G.: An overview of the MOP runtime verification framework. J. Softw. Tools Technol. Transf. **14**, 249–289 (2011)
58. Roşu, G., Chen, F.: Semantics and algorithms for parametric monitoring. Log. Methods Comput. Sci. **8**(1), 1–39 (2012)
59. Roşu, G., Chen, F., Ball, T.: Synthesizing monitors for safety properties: this time with calls and returns. In: Leucker, M. (ed.) RV 2008. LNCS, vol. 5289, pp. 51–68. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-89247-2_4
60. Roşu, G., Havelund, K.: Rewriting-based techniques for runtime verification. Autom. Softw. Eng. **12**(2), 151–197 (2005)
61. Savage, S., Burrows, M., Nelson, G., Sobalvarro, P., Anderson, T.: Eraser: a dynamic data race detector for multithreaded programs. ACM Trans. Comput. Syst. **15**(4), 391–411 (1997)
62. Seyster, J., et al.: InterAspect: aspect-oriented instrumentation with GCC. Form. Methods Syst. Des. **41**(3), 295–320 (2012)
63. Smith, D.R., Havelund, K.: Toward automated enforcement of error-handling policies. Technical Report number: TR-KT-0508, Kestrel Technology LLC, August 2005

64. Stoller, S.D., et al.: Runtime verification with state estimation. In: Khurshid, S., Sen, K. (eds.) RV 2011. LNCS, vol. 7186, pp. 193–207. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-29860-8_15
65. Walker, R., Viggers, K.: Implementing protocols via declarative event patterns. In: Taylor, R., Dwyer, M., (eds.) ACM Sigsoft 12th International Symposium on Foundations of Software Engineering (FSE-12), pp. 159–169. ACM Press (2004)