

Software that Meets Its Intent

Marieke Huisman¹(✉), Herbert Bos², Sjaak Brinkkemper³, Arie van Deursen⁴,
Jan Friso Grooten⁵, Patricia Lago², Jaco van de Pol¹, and Eelco Visser⁴

¹ University of Twente, Enschede, The Netherlands
M.Huisman@utwente.nl

² Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

³ Utrecht University, Utrecht, The Netherlands

⁴ Delft University of Technology, Delft, The Netherlands

⁵ Eindhoven University of Technology, Eindhoven, The Netherlands

Abstract. Software is widely used, and society increasingly depends on its reliability. However, software has become so complex and it evolves so quickly that we fail to keep it under control. Therefore, we propose *intents*: fundamental laws that capture a software systems' intended behavior (resilient, secure, safe, sustainable, etc.). The realization of this idea requires novel theories, algorithms, tools, and techniques to discover, express, verify, and evolve software intents. Thus, future software systems will be able to verify themselves that they meet their intents. Moreover, they will be able to respond to deviations from intents through self-correction. In this article we propose a research agenda, outlining which novel theories, algorithms and tools are required.

1 Introduction

Problem. Software is everywhere: the estimated number of connected devices today ranges in the tens of billions and will be up in the hundreds of billions very soon. While society increasingly depends on its reliability and trustworthiness, the inconvenient truth is that software is very fragile. Software is so complicated that we have no clear idea of what a program may or may not do. And when something is working, we are afraid to touch it. For instance, many critical infrastructures do not update their software for fear of breaking things, even in the case of crucial security updates. As it stands we lack real control over software in virtually all areas. Both criminals and intelligence services employ this to hack into our computers. This is today. Tomorrow will be worse.

Solution. In 1942, the science fiction writer Isaac Asimov formulated a number of laws for robots. For instance, no matter how complex they became, they should never injure a human being, or through inaction, allow a human being to come to harm. In other words, the laws should govern the robot's behavior, independent of its normal functionality or complexity.

Today, some 75 years later, we take up this idea and apply it to software: by defining simple laws or "*intents*" for software, we want to govern a program's

behavior, independent of the implementation of its functionality: it should not crash, it should resist attacks, it should adhere to its functional specification. It should revert buggy updates, it should keep sensitive data private, etc.

Moreover, the intents should be met under any condition – even if the software evolves, or the environment changes. Where needed, the software should correct itself automatically to fulfil the intents. And whenever the intents change, the software has to adjust itself. In case conflicting intents for software are discovered, ethical considerations should be made to resolve the conflict.

This requires a research agenda, to investigate methods for discovering the dynamic intents of living software, techniques to specify and verify intents, and new mechanisms to make software self-correcting, *i.e.*, to make it adapt itself to deviations from the current intents. This requires the concerted efforts from a broad range of software technologies, combining formal analysis, experimental software laboratory research, and empirical research on deployed software. Developing and evaluating the methods also requires intimate collaboration with a number of embedded software application domains, like health care, robotics, automotive, and science experiments.

Software is becoming so versatile, all-embracing and intelligent that it is inconceivable that we can control all detailed aspects of its behavior. But by using intents we can guarantee that software adheres always to its essential purpose, avoids unnecessary harm, and operates within the boundaries of the law. With the same aim as Asimov's robot laws, intents guarantee that software can be used safely and reliably.

2 The Software Reliability Challenge

Software runs the world. Software plays an integral role in our daily lives. There is an 'app' for everything, from the energy grid that powers our infrastructure to the social media that mediate our social interactions. It is hard to recall what life was before all that software, and it is even harder to imagine what we would do without. Society would grind to a halt if the software running it would break down. This dependence on software *will keep growing* in the years to come. Gartner Inc. estimates that the number of connected devices today ranges in the tens of billions and that this will grow to hundreds of billions very soon [42].

The enormous growth in software [87] has led to many advances. Office software and information systems have revolutionized information processing in all sectors, from finance to retail to government. Software is now a crucial and often complex component in many *embedded systems* such as medical surgery robots; control systems for aviation, traffic, and nuclear equipment; and robot companions for elderly people. Software is also having a transforming impact on *science*: scientific experiments are often replaced by much faster, cheaper and safer software simulations, and scientific discoveries often result from automated analysis of large amounts of measurement data.

These successes stem from the fact that *computers outperform humans* in many tasks: computers are faster and more reliable; they can perform many

tasks simultaneously, and can process astounding quantities of data. Already, there are many cases where software ‘knows’ more about the tasks at hand than the humans involved, and this development has only just started.

However, the tremendous growth of software also has a downside. All software systems contain errors that cause them to behave in unintended ways [41, 62]. Studies show that even tested and deployed software has on average between 1 and 16 errors per 1000 lines of code [70, 71]. As a consequence, dormant software faults persist even in mature code bases [20]. With code bases that often exceed a million lines of code, modern software systems easily contain thousands of bugs, many of which may bring down the entire system, or worse, produce wrong results that go undetected. Most software errors are introduced by developers accidentally violating the (informal) expectations on a system. However, errors may also be introduced consciously to actively fool the users [88], which makes error detection even more challenging. Almost daily, we are exposed to news reports on the consequences of failures and other malbehavior of software.

Thus, even though many people think they can depend on the *reliability* of software, the truth is that is just an idealized dream. Software has become *too complex*, and it *evolves* so quickly, that we are not really in control. Using artificial intelligence and autonomous learning techniques, software can adapt itself to its environment to outwit people, operating in a better way than any human could ever have programmed. This makes it scientifically very challenging to predict all possible emergent behaviors of a large, complex, and evolving software system, and thus to ensure that it is and will remain free of flaws, or to guarantee that it will always avoid foul play in obtaining its objectives.

Meanwhile, engineers keep developing new systems, for which new software has to be developed, while the software of old systems grows due to maintenance and adding new features. The resulting complexity is abused by various parties, such as criminals and intelligence services, actively searching for flaws in software systems in order to exploit them for their own advantage. There are even cases where software developers designed “features” in order to mislead customers and users, e.g. to circumvent environmental emission restrictions in cars. As a result, as a society we cannot guarantee the security and integrity of data, ultimately threatening the identity of our citizens. Moreover, as devices become increasingly *connected* and exchange information, the privacy of their users is at stake.

The root cause of this software reliability crisis is that the policies that software systems implement are buried in low-level code, which makes it hard to establish what exactly those policies are, whether they are implemented correctly, and whether these are indeed the policies that we intended. This problem is exacerbated by the rapidly increasing complexity of software systems. This would be akin to a society consisting of individuals that operate by doing as they please, driven by their own particular desires.

By contrast, modern society operates by *rule of law*, imposing constraints on the behavior of its citizens and institutions. While a modern society is a highly complex system with millions of entities interacting in many ways, rule of law keeps this system in check by imposing clear and objective criteria for judging the legality of the behavior of all these entities. To ensure that we can keep

depending on software and to prevent it from running wild, software should be governed by rule of law, just like our society!

3 Software Intents

Software laws, which we will call *software intents*, explicitly describe the essential limits on and expectations on the *behavior* of a software system. Software intents often cover functional requirements, but they should also cover non-functional requirements (e.g., intended performance, response time), and security properties (e.g., resilience against attackers). A crucial property of software intents is that they describe high-level requirements that reflect the goals of a software system, not (just) the low-level effects of a particular realization of those goals on a particular computing substrate.

Common examples of software intents are: Software should not crash; When software crashes nevertheless, it should do so gracefully; Software should not cause any harm to its users, or to any other human being; Software should be resistant to attacks; Software should satisfy its specifications; and: Software should not violate the privacy of its users. For a telling example, consider the recent case of the Japanese satellite Hitomi, which tumbled out of control five weeks after launch, because of a software error [94]: in an attempt to stop itself from spinning, the control system fired a thruster jet in the wrong direction, causing an acceleration, instead of a slow down. An implicit intent for a satellite is that it should stay in its orbit. By making this intent explicit, and providing support to ensure that this intent is fulfilled at all times, the error could have been detected and corrected.

Software intents serve multiple purposes. First, each software system should explicitly *declare* its intents, which serves as a promise and explanation what it intends to do. Second, intents should be usable to *verify* that a software system is indeed realizing its intended behavior (and no others). Finally, intents should facilitate the *evolution* of software systems, by ensuring that software maintenance is intent preserving. Technically, *intents* can be viewed as a form of redundancy, which will allow to increase software reliability fundamentally, *i.e.*, reduce the chance of software errors to less than 10^{-10} .

Therefore, we propose a novel research agenda, revolving around the *paradigm of software intents*. This agenda should address at least the following areas:

- *Expression of Intent*: How to discover and define software intents?
- *Verification of Intent*: How to ensure that software systems *always* behave as intended?
- *Evolution of Intent*: How to ensure that software systems keep behaving as intended, also in the future?

While continuous change, reliability and trustworthiness regard any type of software system, they are especially critical in embedded software where intents derived from safety, ethics and performance concerns play an essential role. Therefore, this research agenda can only be effective in close collaboration with

a selected number of embedded software application domains. Clear candidate domains are robotics, in industry, autonomous automotive, or home care situations, other health care applications, and science experiments in physics and biology. Here, reliability, safety, ethics, and performance play a prominent role.

3.1 Expression of Intent

The first challenge is *how to discover the implicit intents* of software systems. Generally, people have all kinds of expectations on software, but are incapable of making that precise. What are the fundamental properties that should continue to hold throughout the lifetime of a software system, no matter how it evolves? And how can we be sure that all relevant intents have been formulated? Intents might be specified upfront, but another interesting challenge is to discover intents for existing (legacy or open source) systems by observing their use, or mining user or developer data like logs, chats, code repositories, bug trackers, test suites, etc.

A particular interesting question is whether different intents can be *combined* or when intents are *conflicting*. In case of conflicting intents, we need to understand how to resolve these conflicts, which quickly becomes an *ethical question*, as it requires prioritizing one intent over the other. For example, in principle all user data should remain private, but if the user is a suspected terrorist, the privacy restriction might be released in favor of the society's security.

The next challenge is *how to express software intents*. There can be many sources of intents: they can be linguistic expressions of intent by human developers or users; they can be legal restrictions, imposed by laws on safety, privacy and security; or formal specifications, models, test cases etc. Thus, intents may be formal, but they do not have to be. How should suitable *languages of intent*, covering different classes of intents, be defined? For example, as a language of intent for a specific application domain, or as a language of intent for a particular class of properties? Rather than requiring a single language, we stress the usability and understandability of the languages, so that when an intent is expressed, all stakeholders have a common understanding of how the software should behave. Thus, these languages of intent will provide a common framework for all developers and users of software to express what they believe to be the intended program's behavior. Moreover, the advantage of having a well-defined language is that its underlying *semantics of intents* can be defined precisely and formally. This makes it possible to demonstrate that intents are compatible, and ultimately that software satisfies its intent.

3.2 Verification of Intent

There are many different ways to verify intents on software. This can be done at *design time* using logic-based techniques, which can provide high-correctness guarantees. It is possible to do this by testing applications in a *software laboratory* setting before deployment. But it can also be done after deployment, by *monitoring applications* in real-life situations in the specific environment in which the software is run.

Design-time verification (or static verification), is done on the basis of models of the program code or on the code itself. Instead of executing the program, it analyses all the possible program executions (for example, by exploring the complete state space of a program, or by representing program executions at a suitably abstract level in a symbolic way). Over the last decade, significant progress has been made to apply design-time verification to large-scale examples [27, 53, 69]. However design-time verification still requires a significant amount of user intervention. For design-time verification of intent-based software, the major research challenges that need to be addressed are:

1. to *fully automate* this process by automatically generating all the information that is necessary to capture all possible program behaviors in a suitable way;
2. to investigate how design-time verification can be *generalized* to other classes of intents than traditional functional specifications; and
3. to apply design-time verification when one only has control over parts of the software system only.

Testing, or run-time verification, considers concrete program executions, and verifies in a software laboratory whether software does not violate the intended behavior of the program. Because of the lab situation, one has full control over the program, and can rerun it, applying smart (or dangerous) test cases. The main challenges are:

1. to guide the testing process to test cases that might lead to discovering potential intent violations, or can support certifying their absence;
2. while at the same time optimizing the time and effort put in the testing process, possibly taking into account a suitable risk assessment.

Finally, **monitoring** is needed to verify that the software fulfils its intent even after software deployment. Verification and testing typically assume a particular environment in which the software is run, but after deployment the environment might change in unforeseen ways. Monitoring provides a first step, when a violation of intent is detected or predicted, to adapt the behavior of the system, or if necessary to completely interrupt it before any real harm can be done. The major challenges are to understand how the monitoring can be done effectively for all classes of intents:

1. without significant *overhead on the performance* of the software; and
2. in the presence of *great uncertainty* due to the dynamic environment in which the software is running.

3.3 Evolution of Intent

A fact of life (which makes it interesting) is that both the intents that software should satisfy as well as the environment in which it must run are continuously changing. If the environment changes, the software should live up to its intents, and if needed, adjust itself automatically. In case the intent itself changes, for

example because of a change in user requirements, or a new law imposing new rules that the software must satisfy, the software has to be adjusted to live up to its intents. In traditional software engineering, the software has to be re-engineered. Because of the enormous size of software this is a costly, error prone and time consuming affair. Avoiding this requires novel techniques such that the software can correct itself autonomously, and then provide guarantees that it behaves according to its (new) intent again. Exciting new research shows that it is possible to achieve this on small programs, with search-based techniques, [58, 73]. However, turning this into a scalable technique will require a major effort.

The first challenge is to investigate how software can *establish* itself that it is still fulfilling its intents, especially when there are changes in the environment. The software *should take responsibility for its own intent*. The next challenge is to understand the different techniques for *efficient correction* of the software in case the intents are violated. Due to the size and complexity of software, repairing software at one place can have an unwanted and unexpected impact on other parts of its behavior. Efficient methods are required, to show that local software reparations do not adversely affect the rest of the code. The challenge is to show that the intent is fulfilled again without a complete re-verification. Finally, the last major challenge is to understand if and how these correction techniques can be applied autonomously, i.e., to make the software *self-correcting*. For this, techniques are required that select the appropriate corrections for the diagnosed source of the problem automatically, and then apply these corrections.

Just as modern spellcheckers effectively apply auto-correction to natural language texts by pattern recognition, in our vision future software systems should

1. *detect software patterns* that violate their own intents; and
2. *identify and apply correction patterns automatically* to make software respect its intents again.

4 Current State-of-the-Art

The quest for reliable software is not new, but started right after the conception of programmable machines. In this section, we provide a broad overview of the current state of the art on the topics covered by the proposed research agenda.

4.1 Expression of Intent

Requirements engineering is traditionally used to understand and model the software's intended behavior. Requirements engineering typically focuses on identifying, modeling and analyzing the goals of the stakeholders [89]. Usually, every stakeholder has a different perspective on what the software should do [37, 78]. These *viewpoints* can be independent or partially overlapping. Various proposals exist to resolve the inconsistencies between viewpoints [5, 36]. A common approach is to prioritize the requirements.

A systematic approach to incorporate ethical and legal questions into requirement prioritization is still largely missing. From a legal perspective, the formulation of norms and regulations on software is slowly gaining traction, especially in the field of security and privacy [15, 54]. Recently, [24] formalized intents about a wide range of security requirements. Lago et al. started a library of intents in the domain of cloud-based software [61, 76]. Despite their long history, the adoption of goal modelling techniques is still low. Software requirements are often written in natural language [90], so they cannot be linked directly to the other activities of software development.

At the other side of the spectrum, we have formal languages. For instance, modal logics specify external system behaviour [8, 48]. They operate on idealized system models rather than on actual code. For program verification, several program annotation languages have been introduced [51, 59]. These approaches are precise, but limited in expressivity. They are based on formal semantics, detailing the intended behavior of substantial parts of widely used programming languages, like C [34, 60] and sequential and concurrent Java [6, 52]. Other formal specification notations, like B and VDM, operate at the system level [3, 38].

Formal notations are precise and often very expressive, but hard to use, even by insiders. Domain-specific languages (DSLs) [40, 86] provide a good middle ground between informal natural language and low-level formal logics. DSLs provide domain-specific notation, analysis, verification, and optimization. A key complication is the cost of their implementation and maintenance, as studied by Van Deursen et al. [84]. This can be alleviated by the development of programming environments for the domain-specific programming languages [35, 39]. A limitation of the current generation of language workbenches is the lack of support for guaranteeing the safety and reliability of languages, which is crucial for the software intent paradigm. This requires integrating verification of language properties in language workbenches [68].

4.2 Verification of Intent

In industry, testing is currently the prevalent way to discover faults, and to assess the quality of software. It became the most expensive part of system development, calling for test automation to reduce the costs. However, testing is poorly understood: while testing does reveal many software faults, it is currently unknown how much testing is sufficient to gain some justified level of confidence. Also, testing is often focused at small components (unit testing) while the actual challenges arise at system integration. An advanced test automation technique for control-intensive software is model-based testing [82]. This provides a sound and complete theory for deriving, executing, and evaluating test cases from a model of a system. However, systematic research on computing the next best optimal test case is still missing. For data generation, search-based testing has been investigated [79], which uses meta-heuristic search for steering random test suites towards optimizing a test fitness function. For both technologies, there is a gap between the intents of the system and the used test models.

Another challenge is to establish the likelihood that software satisfies its intent after testing, or to generate optimal test cases that lower that likelihood. This challenge is tough. An overview of the field is presented in [96], providing nice and intricate theoretical testing models. Aligning these models with practical observation is a big challenge. Viewing testing as a learning activity is promising, but not yet widespread. It is most prominent in the field of *exploratory testing* [50, 93]. Our notion of intent could be used to record the knowledge obtained by exploratory testing. An algorithmic interpretation is provided by automaton learning [19, 30, 63, 64], which has been used complementary to testing [1].

Automated verification of software at design-time raises confidence in program correctness. Its two fundamental pillars are modularity and automated reasoning. *Modularity* forms the cornerstone to scale up verification technology to full systems. It is rooted in proper abstractions of program behavior. In his seminal work, Robin Milner [66] introduced bisimulation as the appropriate equivalence notion. This allowed to reason about the abstract behavior of a program in its wider context. Hoare logic [51] provides modularity at the level of program code. Pre- and post-conditions abstract away from the implementation of software functions. Other parts of the program need to deal with their specification only.

Automated reasoning has been revolutionized by SAT solving. Although SAT solving is NP-hard, a series of breakthroughs [17, 25, 33, 49, 80] led to algorithms that allowed to solve instances of propositional formulas of tremendous size. Many problems have been translated into SAT solving. An early example is the verification of PLC programs by Groote [47] that guard the safety of railway yards, bridges and sluices. More recently, SAT solvers were extended with mathematical theories, leading to SMT solvers (satisfiability modulo theories), like Z3 and Mathsat [21, 29], reasoning about data types automatically.

The three main branches in automated program verification are: (1) *Interactive proof checkers*. The Dutch mathematician De Bruijn observed in 1967 that proof checking could be mechanized. He built Automath [26], the first program to check mathematical proofs. Modern proof checkers, like Coq [11], follow the same principle. The user constructs a proof interactively, which is checked by an independent procedure, to eliminate all reasonable doubt on the validity of a theorem. This approach is versatile, but not feasible for non-trivial software, since a correctness proof is extremely detailed.

(2) *Model checkers* [8], invented by the Turing Award winners Clarke, Emerson and Sifakis, operate fully automatically in principle. They take a (model of a) program and a specification, and check the property by brute force computation. One of the most expressive model checkers is the mCRL2 toolset [48]. Tremendous progress was made by automated abstraction techniques [14, 22], the use of BDD and SAT-solving technology [12], and the use of high-performance graph algorithms for model checking, e.g. [55]. Model checkers operate on finite systems of limited size only; checking realistic software with datatypes and concurrency remains a challenge.

(3) *Automated program verifiers*, like OpenJML and KeY [4,23] strike a balance between expressivity and automation. They directly operate on software programs, generating a sufficient set of verification conditions, which are automatically discharged by SMT solvers. Non-trivial programs can be handled, which is demonstrated by revealing an error in the popular TIMsort routine in the Java library with KeY [27]. Extending Hoare Logic [51] by Separation Logic [32] allowed handle concurrency and data structures of modern software. This led to tools that can reason about parallel Java programs, like the VerCors tool [13]. Still, program verifiers can only be used by verification experts, since they require manual annotations of pre- and post-conditions, as well as auxiliary properties like invariants and thread permissions. This doesn't scale beyond medium size programs.

Monitoring deployed software is common to all programming languages and even computer hardware. An example is checking memory accesses, to prevent violation of the intent that programs should access only their own memory. This can be supported by hardware, guaranteed by the Operating System, and is built-in in many modern programming languages (like Java and C#).

In the security context, the versatility and unpredictability of attacks on software make it hard to verify against *a priori* defined intent. Besides monitoring violation of well-defined security intents, software is often monitored globally and any unusual activity (such as extreme network activity) is reported for further investigation. Whenever a security breach is detected, this is generally transformed into an explicit intent that can be checked. Advanced security monitoring techniques are pioneered by Bos et al [81,83]. Likewise, Lago's team has pioneered monitoring of sustainability intents [57,75,77].

In the run-time verification community, attention has shifted to the dynamic verification of behavioural properties on programs [10,31]. These techniques are often limited to safety properties, which guarantee absence of malbehavior. Like user-inserted asserts, all these forms of monitoring typically lead to a stop-the-world whenever some intent violation is detected.

4.3 Evolution of Intent

Programming by contract [65] emerged as a systematic approach. When the software intent or the environment change, it can be checked whether the environment still offers sufficient functionality, such that the software can still guarantee its renewed contract. To establish that the software itself ends up undamaged in core memory, traditional methods are self-tests, checksums and oblivious hashing. These can be circumvented, so security researchers introduced tamper-proof software, which relies on encryption [7,95].

To localize faults, spectrum-based fault diagnosis attributes observed failures to the responsible software components [2,74]. A different perspective is origin tracking [85], where the contribution of a certain source element to some external effect is carefully traced. A promising new localisation approach is proof carrying code [9,67]. Here, the program correctness proof is encoded within the software.

If the intent changes, proof reconstruction might point to the relevant parts of the proof and program that must be changed.

There are a number of fields where returning to a safe situation has been investigated in depth. Examples are file storage [91], database transactions [56] and networking [18]. Generally, such techniques rely on external operators to restart systems and bring them back in a working state. A whole research line on software engineering considers self-* systems (e.g. self-adaptation, self-management, self-healing, self-configuration). Here a system must preserve its operation at run-time even in the presence of changes in environment, resources, security breaches, faults [28]. Solutions come from various areas, like software architecture, fault-tolerance, and bio-inspired computing [16, 72, 92]. These techniques typically protect against environment uncertainty.

If the software contains errors, or when the intents themselves change, software must be updated. There are several techniques to update software while it is running (hot swappable). Reliable dynamic updates and error recovery are studied in [44, 45]. There is little systematic research into the question how software can handle its own mishap completely autonomously. Only very recently, some indications that self-correcting code might work were published [43, 58, 73]. The question is whether these observations carry over to software on a larger scale. In particular, it will be challenging to compute the best update on the software that removes the problem.

The fundamental challenge in self-correcting code lies in guaranteeing that code adapts itself in such a way that it will continue to behave in the right way, even if the intent has not been formulated explicitly and completely. In the past, various forms of self-adaptation of software have been used, generally with an adverse effect. An excellent example is the programming language COMAL [46], which would correct simple programming errors by itself. Superficial typos were sometimes transformed in deep, hard to spot programming errors, undoing the beneficiary effect of type checking.

5 Perspectives

We conclude by summarizing our perspective on achieving software that meets its intents. In particular, we identify technical challenges that will need further research in the next decade.

First, intents must be discovered and expressed. The goal is that all stakeholders can express, understand, and prioritize their expectations on software. Domain specific languages should bridge the enormous gap between human intent and low-level machine readable code.

Second, intents must be verifiable, both statically during design (verification), dynamically in a laboratory setting (testing), and on deployed code (monitoring). The big challenge is to make an automation leap: human effort to annotate or instrument programs, construct proofs, derive test cases, or operate verification tools should be avoided. Technically, this requires an advance of annotation and invariant generation. At the same time, verification should bridge the gap between detailed code and its high-level, emerging properties.

Finally, software should correct itself when it contains errors, when its environment changes, or when the intents of people or organizations evolve. Current runtime adaptation and software update technology should be complemented by an auto-correction facility for code. Building in explicit and executable notions of intent make it possible in principle that software adapts itself to its intent.

These technical developments require the coordinated effort of researchers in theoretical computer science, programming languages, and empirical software engineering. When successful, they could form the basis for intents as software laws that can be verified and guaranteed. Only automatically self-correcting software will enable us to regain grip on the complex, evolving software systems that run our society. Ultimately, this development allows that software providers can be held accountable for their product.

References

1. Aarts, F., Kuppens, H., Tretmans, J., Vaandrager, F.W., Verwer, S.: Improving active Mealy machine learning for protocol conformance testing. *Mach. Learn.* **96**(1–2), 189–224 (2014)
2. Abreu, R., Zoetewij, P., Van Gemund, A.J.C.: A new Bayesian approach to multiple intermittent fault diagnosis. In: *International Joint Conference on Artificial Intelligence, IJCAI 2009*, pp. 653–658 (2009)
3. Abrial, J.: *Modeling in Event-B - System and Software Engineering*. Cambridge University Press, Cambridge (2010)
4. Ahrendt, W., Beckert, B., Hähnle, R., Rümmer, P., Schmitt, P.H.: Verifying object-oriented programs with KeY: a tutorial. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) *FMCO 2006*. LNCS, vol. 4709, pp. 70–101. Springer, Heidelberg (2007)
5. Ali, R., Dalpiaz, F., Giorgini, P.: Reasoning with contextual requirements: detecting inconsistency and conflicts. *Inf. Softw. Technol.* **55**(1), 35–57 (2013)
6. Amighi, A., Haack, C., Huisman, M., Hurlin, C.: Permission-based separation logic for multithreaded Java programs. *Logical Methods Comput. Sci.* 11(1:2), 1–66 (2015) paper 2
7. Andriess, D., Bos, H., Slowinska, A.: Parallax: implicit code integrity verification using return-oriented programming. In: *IEEE/IFIP IC on Dependable Systems and Networks, DSN 2015*, pp. 125–135. IEEE Computer Society (2015)
8. Baier, C., Katoen, J.-P.: *Principles of Model Checking*. The MIT Press, Cambridge (2008)
9. Barthe, G., Grégoire, B., Kunz, C., Rezk, T.: Certificate translation for optimizing compilers. *ACM Trans. Program. Lang. Syst.* **31**(5), 18 (2009)
10. Bauer, A., Leucker, M., Schallhart, C.: Runtime verification for LTL and TLTL. *ACM Trans. Softw. Eng. Methodol.* **20**(4), 14 (2011)
11. Bertot, Y., Castéran, P.: *Interactive Theorem Proving and Program Development - Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. An EATCS Series. Springer, Heidelberg (2004)
12. Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. *Adv. Comput.* **58**, 117–148 (2003)
13. Blom, S., Huisman, M.: The VerCors tool for verification of concurrent programs. In: Jones, C., Pihlajasaari, P., Sun, J. (eds.) *FM 2014*. LNCS, vol. 8442, pp. 127–131. Springer, Heidelberg (2014)

14. Bradley, A.R.: IC3 and beyond: incremental, inductive verification. In: Madhusudan, P., Seshia, S.A. (eds.) CAV 2012. LNCS, vol. 7358, p. 4. Springer, Heidelberg (2012)
15. Breaux, T.D., Vail, M.W., Antón, A.I.: Towards regulatory compliance: extracting rights and obligations to align requirements with regulations. In: IEEE International Requirements Engineering Conference, pp. 46–55 (2006)
16. Brun, Y., Di Marzo Serugendo, G., Gacek, C., Giese, H., Kienle, H., Litoiu, M., Müller, H., Pezzè, M., Shaw, M.: Engineering self-adaptive systems through feedback loops. In: Cheng, B.H.C., Lemos, R., Giese, H., Inverardi, P., Magee, J. (eds.) Software Engineering for Self-Adaptive Systems. LNCS, vol. 5525, pp. 48–70. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-02161-9_3](https://doi.org/10.1007/978-3-642-02161-9_3)
17. Bryant, R.E.: Symbolic manipulation of Boolean functions using a graphical representation. In: Ofek, H., O’Neill, L.A. (eds.) 22nd ACM/IEEE Conference on Design Automation, (DAC 1985), pp. 688–694. ACM (1985)
18. Cachin, C., Guerraoui, R., Rodrigues, L.E.T.: Introduction to Reliable and Secure Distributed Programming, 2nd edn. Springer, Heidelberg (2011)
19. Cassel, S., Howar, F., Jonsson, B., Steffen, B.: Active learning for extended finite state machines. *Formal Aspects Comput.* **28**(2), 233–263 (2016)
20. Chen, T.-H., Nagappan, M., Shihab, E., Hassan, A.E.: An empirical study of dormant bugs. In: 11th Working Conference on Mining Software Repositories, MSR 2014, pp. 82–91. ACM (2014)
21. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The MathSAT5 SMT solver. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013 (ETAPS 2013). LNCS, vol. 7795, pp. 93–107. Springer, Heidelberg (2013)
22. Clarke, E.M., Gupta, A., Strichman, O.: SAT-based counterexample-guided abstraction refinement. *IEEE Trans. CAD Integr. Circ. Syst.* **23**(7), 1113–1123 (2004)
23. Cok, D.R.: OpenJML: software verification for Java 7 using JML, OpenJDK, and Eclipse. In: Dubois, C., Giannakopoulou, D., Méry, D. (eds.) 1st Workshop on Formal Integrated Development Environment, (F-IDE 2014). EPTCS, vol. 149, pp. 79–92 (2014)
24. Dalpiaz, F., Paja, E., Giorgini, P.: Security Requirements Engineering: Designing Secure Socio-Technical Systems, 1st edn. MIT Press, Cambridge (2016)
25. Davis, M., Logemann, G., Loveland, D.W.: A machine program for theorem-proving. *Commun. ACM* **5**(7), 394–397 (1962)
26. de Bruijn, N.: A survey of the project AUTOMATH. In: To H.B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism, pp. 579–606. Academic Press (1980)
27. de Gouw, S., Rot, J., de Boer, F.S., Bubel, R., Hähnle, R.: OpenJDK’s Java.util.Collection.sort() is broken: the good, the bad and the worst case. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 273–289. Springer, Heidelberg (2015)
28. de Lemos, R., et al.: Software engineering processes for self-adaptive systems. In: Lemos, R., Giese, H., Müller, H.A., Shaw, M. (eds.) Software Engineering for Self-Adaptive Systems II. LNCS, vol. 7475, pp. 51–75. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-35813-5_3](https://doi.org/10.1007/978-3-642-35813-5_3)
29. de Moura, L., Bjørner, N.S.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
30. van Deursen, A.: Testing web applications with state objects. *Commun. ACM* **58**(8), 36–43 (2015)

31. Diekert, V., Leucker, M.: Topology, monitorable properties and runtime verification. *Theor. Comput. Sci.* **537**, 29–41 (2014)
32. Distefano, D., O’Hearn, P.W., Yang, H.: A local shape analysis based on separation logic. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920, pp. 287–302. Springer, Heidelberg (2006)
33. Eén, N., Sörensson, N.: An extensible SAT-solver. In: Giunchiglia, E., Tacchella, A. (eds.) SAT 2003. LNCS, vol. 2919, pp. 502–518. Springer, Heidelberg (2004)
34. Ellison, C., Rosu, G.: An executable formal semantics of C with applications. In: 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL 2012), pp. 533–544. ACM (2012)
35. Erdweg, S., et al.: Evaluating and comparing language workbenches: existing results and benchmarks for the future. *Comput. Lang. Syst. Struct.* **44**, 24–47 (2015)
36. Finkelstein, A., Gabbay, D., Hunter, A., Kramer, J., Nuseibeh, B.: Inconsistency handling in multiperspective specifications. *IEEE TSE* **20**(8), 569–578 (1994)
37. Finkelstein, A., Kramer, J., Nuseibeh, B., Finkelstein, L., Goedicke, M.: Viewpoints: a framework for integrating multiple perspectives in system development. *Int. J. Softw. Eng. Knowl. Eng.* **2**(1), 31–57 (1992)
38. Fitzgerald, J., Larsen, P.G.: *Modelling Systems: Practical Tools and Techniques for Software Development*, 2nd edn. Cambridge University Press, Cambridge (2009)
39. Fowler, M.: Language workbenches: The killer-app. for domain specific languages? (2005). <http://www.martinfowler.com/articles/languageWorkbench.html>
40. Fowler, M.: *Domain-Specific Languages*. Addison Wesley, Boston (2010)
41. Ganapathi, A., Patterson, D.A.: Crash data collection: a windows case study. In: DSN, pp. 280–285. IEEE Computer Society (2005)
42. Gartner Inc. Smart cities will include 10 billion things by 2020 (2015). <https://www.gartner.com/doc/3004417/smart-cities-include-billion>
43. Ghardallou, W., Diallo, N., Mili, A.: Program derivation by correctness enhancements. In: *Refinement* (2015)
44. Giuffrida, C., Cavallaro, L., Tanenbaum, A.S.: Practical automated vulnerability monitoring using program state invariants. In: DSN, October 2013
45. Giuffrida, C., Iorgulescu, C., Kuijsten, A., Tanenbaum, A.S.: Back to the future: fault-tolerant live update with time-traveling state transfer. In: LISA, October 2013
46. Gratte, I.: *Starting with COMAL*. Prentice-Hall, Englewood Cliffs (1985)
47. Groote, J., Koorn, J., van Vlijmen, S.: The safety guaranteeing system at station Hoorn-Kersenboogerd (extended abstract). In: 10th Annual Conference on Computer Assurance (COMPASS 1995), pp. 57–68 (1995)
48. Groote, J., Mousavi, M.: *Modeling and Analysis of Communicating Systems*. The MIT Press, Cambridge (2014)
49. Groote, J.F., Warners, J.P.: The propositional formula checker HeerHugo. *J. Autom. Reasoning* **24**(1/2), 101–125 (2000)
50. Hendrickson, E.: *Explore It!: Reduce Risk and Increase Confidence with Exploratory Testing*. The Pragmatic Bookshelf, Raleigh (2013)
51. Hoare, C.: An axiomatic basis for computer programming. *Commun. ACM* **12**(10), 576–580 (1969)
52. Huisman, M.: *Reasoning about Java Programs in Higher Order Logic with PVS and Isabelle*. Ph.D. thesis, University of Nijmegen (2001)
53. Hwong, Y., Keiren, J., Kusters, V., Leemans, S., Willemse, T.: Formalising and analysing the control software of the compact muon solenoid experiment at the large hadron collider. *Sci. Comput. Program.* **78**, 2435–2452 (2013)

54. Ingolfo, S., Siena, A., Mylopoulos, J.: Establishing regulatory compliance for software requirements. In: Jeusfeld, M., Delcambre, L., Ling, T.-W. (eds.) ER 2011. LNCS, vol. 6998, pp. 47–61. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-24606-7_5](https://doi.org/10.1007/978-3-642-24606-7_5)
55. Kant, G., Laarman, A., Meijer, J., van de Pol, J., Blom, S., van Dijk, T.: LTSmin: high-performance language-independent model checking. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 692–707. Springer, Heidelberg (2015)
56. Kemme, B., Jiménez, R., Patiño-Martínez, M.: Database Replication. Synthesis Lectures on Data Management. Morgan & Claypool Publishers, San Rafael (2010)
57. Lago, P., Koçak, S.A., Crnkovic, I., Penzenstadler, B.: Framing sustainability as a property of software quality. *Commun. ACM* **58**(10), 70–78 (2015)
58. Le Goues, C., Nguyen, T., Forrest, S., Weimer, W.: GenProg: a generic method for automatic software repair. *IEEE Trans. Softw. Eng.* **38**(1), 54–72 (2012)
59. Leavens, G., Poll, E., Clifton, C., Cheon, Y., Ruby, C., Cok, D.R., Müller, P., Kiniry, J., Chalin, P.: JML Reference Manual, Dept. of Computer Science, Iowa State University, February 2007. <http://www.jmlspecs.org>
60. Leroy, X.: Formal verification of a realistic compiler. *Commun. ACM* **52**, 107–115 (2009)
61. Lewis, G.A., Lago, P., Avgeriou, P.: A decision model for cyber-foraging systems. In: Proceedings of the 13th Working IEEE/IFIP Conference on Software Architecture (WICSA 2016), pp. 51–60. IEEE (2016)
62. Matias, R., Prince, M., Borges, L., Sousa, C., Henrique, L.: An empirical exploratory study on operating system reliability. In: 29th Annual ACM Symposium on Applied Computing, SAC 2014, pp. 1523–1528. ACM (2014)
63. Mesbah, A., van Deursen, A., Lenselink, S.: Crawling Ajax-based web applications through dynamic analysis of user interface state changes. *ACM Trans. Web* **6**(1), 3 (2012)
64. Mesbah, A., van Deursen, A., Roest, D.: Invariant-based automated testing of modern web applications. *IEEE Trans. Softw. Eng.* **38**(1), 35–53 (2012)
65. Meyer, B.: Touch of Class: Learning to Program Well with Objects and Contracts. Springer, Heidelberg (2009)
66. Milner, R.: Calculus of Communicating Systems. Lectures in Computer Science, vol. 92. Springer, Heidelberg (1980)
67. Necula, G.C.: Proof-carrying code. In: Principles of Programming Languages (1997)
68. Neron, P., Tolmach, A., Visser, E., Wachsmuth, G.: A theory of name resolution. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 205–231. Springer, Heidelberg (2015)
69. Osaiweran, A., Schuts, M., Hooman, J., Groote, J., van Rijnsoever, B.: Evaluating the effect a lightweight formal technique in industry. *Int. J. Softw. Tools Technol. Transf.* **18**, 93–108 (2016)
70. Ostrand, T.J., Weyuker, E.J.: The distribution of faults in a large industrial software system. In: 2002 ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2002, pp. 55–64. ACM (2002)
71. Ostrand, T.J., Weyuker, E.J., Bell, R.M.: Where the bugs are. In: 2004 ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2004, pp. 86–96. ACM (2004)
72. Patikirikorala, T., Colman, A., Han, J., Wang, L.: A systematic survey on the design of self-adaptive software systems using control engineering approaches. In: Proceedings of the 7th International Symposium on Software Engineering for Adaptive and Self-Managing Systems, pp. 33–42. IEEE Press (2012)

73. Pei, Y., Furia, C.A., Nordio, M., Wei, Y., Meyer, B., Zeller, A.: Automated fixing of programs with contracts. *IEEE Trans. Softw. Eng.* **40**(5), 427–449 (2014)
74. Perez, A., Abreu, R., van Deursen, A.: A unifying metric for test adequacy and diagnosability. In: *Automated Software Engineering* (2016, Submitted)
75. Procaccianti, G., Fernández, H., Lago, P.: Empirical evaluation of two best practices for energy-efficient software development. *J. Syst. Softw.* **117**, 185–198 (2016)
76. Procaccianti, G., Lago, P., Lewis, G.A.: A catalogue of green architectural tactics for the cloud. In: *Maintenance and Evolution of Service-Oriented and Cloud-Based Systems (MESOCA 2014)*, pp. 29–36. IEEE (2014)
77. Procaccianti, G., Lago, P., Vetro, A., Fernández, D.M., Wieringa, R.: The green lab: experimentation in software energy efficiency. In: *Proceedings of the 37th International Conference on Software Engineering-Volume 2* (2015)
78. Rozanski, N., Woods, E.: *Software Systems Architecture: Working with Stakeholders using Viewpoints and Perspectives*. Addison-Wesley, Boston (2012)
79. Salvesen, K., Galeotti, J.P., Gross, F., Fraser, G., Zeller, A.: Using dynamic symbolic execution to generate inputs in search-based GUI testing. In: Gay, G., Antoniol, G. (eds.) *8th IEEE/ACM International Workshop on Search-Based Software Testing, SBST 2015*, pp. 32–35. IEEE (2015)
80. Sheeran, M., Stålmarck, G.: A tutorial on Stålmarck’s proof procedure for propositional logic. *Formal Methods Syst. Des.* **16**(1), 23–58 (2000)
81. Slowinska, A., Stancescu, T., Bos, H.: Body armor for binaries: preventing buffer overflows without recompilation. In: *Proceedings of USENIX Annual Technical Conference*, Boston, MA, June 2012
82. Tretmans, J.: Model based testing with labelled transition systems. In: Hierons, R.M., Bowen, J.P., Harman, M. (eds.) *FORTEST*. LNCS, vol. 4949, pp. 1–38. Springer, Heidelberg (2008)
83. van der Veen, V., Goktas, E., Contag, M., Pawlowski, A., Chen, X., Rawat, S., Bos, H., Holz, T., Athanasopoulos, E., Giuffrida, C.: A tough call: mitigating advanced code-reuse attacks at the binary level. In: *Proceedings of the 37th IEEE Symposium on Security and Privacy (Oakland)*, San Jose, CA, USA, IEEE, May 2016
84. van Deursen, A., Klint, P.: Little languages: little maintenance? *J. Softw. Maintenance* **10**(2), 75–92 (1998)
85. van Deursen, A., Klint, P., Tip, F.: Origin tracking. *J. Symbolic Comput.* **15**(5/6), 523–545 (1993)
86. van Deursen, A., Klint, P., Visser, J.: Domain-specific languages: an annotated bibliography. *SIGPLAN Not.* **35**(6), 26–36 (2000)
87. van Genuchten, M., Hatton, L.: Metrics with impact. *IEEE Softw.* **30**, 99–101 (2013)
88. van Genuchten, M., Hatton, L.: When software crosses a line. *IEEE Softw.* **33**, 29–31 (2016)
89. van Lamsweerde, A.: Requirements engineering in the year 00: a research perspective. In: *Proceedings of the IEEE International Symposium on Requirements Engineering*, pp. 5–19 (2000)
90. van Lamsweerde, A.: *Requirements Engineering: From System Goals to UML Models to Software Specifications*. Wiley, Hoboken (2009)
91. von Hagen, W., Filesystems, U.: *UNIX Filesystems: Evolution, Design, and Implementation*. SAMS, Indianapolis (2002)
92. Weyns, D., Iftikhar, M.U., de la Iglesia, D.G., Ahmad, T.: A survey of formal methods in self-adaptive systems. In: *Proceedings of the IC on Computer Science and Software Engineering*, pp. 67–79. ACM (2012)

93. Whittaker, J.A.: *Exploratory Software Testing: Tips, Tricks, Tours, and Techniques to Guide Test Design*. Addison-Wesley, Boston (2009)
94. Witze, A.: Software error doomed Japanese Hitomi spacecraft. *Nature* **533**, 18–19 (2016)
95. Wurster, G., van Oorschot, P.C., Somayaji, A.: A generic attack on checksumming-based software tamper resistance. In: *2005 IEEE Symposium on Security and Privacy (S&P 2005)*, pp. 127–138. IEEE Computer Society (2005)
96. Yamada, S.: *Software Reliability Modeling*. Springer, Heidelberg (2014)