

# Correctness-by-Construction and Post-hoc Verification: Friends or Foes?

Maurice H. ter Beek<sup>1</sup>, Reiner Hähnle<sup>2</sup>, and Ina Schaefer<sup>3</sup>

<sup>1</sup> ISTI-CNR, Pisa, Italy

<sup>2</sup> TU Darmstadt, Germany

<sup>3</sup> TU Braunschweig, Germany

**Abstract.** While correctness-by-construction and post-hoc verification are traditionally considered to provide two opposing views on proving software systems to be free from errors, nowadays numerous techniques and application fields witness initiatives that try to integrate elements of both ends of the spectrum. The ultimate aim is not merely to improve the correctness of software systems but also to improve their time-to-market, and to do so at a reasonable cost. This track brings together researchers and practitioners interested in the inherent ‘tension’ that is usually felt when trying to balance the pros and cons of correctness-by-construction versus post-hoc verification.

## Motivation and Aim

Correctness-by-Construction (CbC) sees the development of software (systems) as a scientific discipline of engineering. Originally intended as a mere means of programming algorithms that are correct *by construction* [22,27], the approach found its way into commercial development processes of complex systems [24,25]. In this larger context, we can say that CbC advocates a step-wise refinement process from specification to code, ideally by CbC design tools that automatically generate error-free software implementations from rigorous and unambiguous specifications of requirements. Afterwards, testing only serves the purpose of validating the CbC process rather than to find bugs. (Of course, bugs might still be present outside the boundaries of the verified system: libraries, compilers, hardware, etc.)

In Post-hoc Verification (PhV), on the other hand, formal methods and tools are applied only after the (software) system has been constructed, not during the development process. Typically, a formal specification of (an abstraction of) the implemented system describes how it should behave, after which validation and verification techniques like testing [12,32,33], bug finding [4,26,29], model checking [3,15,17], and deductive verification [7,23,35] are used to check whether the implementation indeed satisfies the specifications and meets the user’s needs. While two independent system models that are verified against each other can provide additional assurance that the designers’ intentions have been captured correctly, PhV is notoriously difficult to carry out.

Recently, numerous techniques and fields of application witness initiatives that attempt to integrate elements of both ends of the spectrum, ultimately aiming to satisfy the holy grail of improving the correctness of software systems as well as their time-to-market, and to do so at a reasonable cost. This track brings together researchers and practitioners interested in the inherent ‘tension’ that is usually felt when trying to balance the pros and cons of CbC vs. PhV. In particular, we invited researchers and practitioners working in the following communities to shed their light on CbC vs. PhV:

- ✓ People working in Software Product Line Engineering (SPLE), who try to lift successful formal methods and verification tools from single product (system) engineering to SPLE in order to say something about the correctness of all products (programs, variants) of an SPL (whose population is exponential in the number of features).
- ✓ People working in System-of-Systems Engineering (SoSE), who address the verification (correctness, but also issues like reliability, resilience, robustness, security, and sustainability) of networks of interacting legacy and new software (systems).
- ✓ People working on (system) synthesis, who aim for transforming a logical specification into a system that is guaranteed to satisfy the specification in all possible environments.
- ✓ People working on deductive verification, who typically require a detailed understanding of why the system works correctly before actual verification to be able to express the correctness of a program as a set of verification conditions to be discharged.
- ✓ People working on ‘bug-finding’ lightweight verification, who trade off full functional verification for being able to deal with real-world languages and large programs as well as to avoid having to write formal specifications.
- ✓ People working on Design for Verification (DfV), mostly in hardware design, who advocate the usage of design methodologies, languages, patterns, etc., that make PhV a realistic option.
- ✓ People working on Statistical Model checking (SMC), who trade off model checking’s verification accuracy, which however requires the entire state space to be known upfront, for scalability by resorting to the computationally more efficient sampling of simulations of (dynamic, black-box, infinite-state) systems until sufficient statistical evidence has been found.

## Contributions

Watson et al. [36] argue for the marriage of CbC with PhV in order to leverage the advantages and to mitigate the disadvantages of both approaches. CbC specifies a problem in terms of its pre- and postconditions and then develops a final algorithmic solution in small, tractable refinement steps. The paper advocates a lightweight approach to proving the correctness of each refinement step. The consequent risk of errors should then be minimised by relying on a PhV system that now obtains *for free* the pre- and postconditions, as well as loop variants and invariants that it requires for proving partial correctness as well as termination.

Beckert et al. [6] address the following important question in the specific setting of legacy code: why is it so hard to perform PhV (deductive verification in particular) and what can be done to make it any easier? They answer the first part of this question by presenting a collection of (known) insights in a systematic way, larded with examples. Subsequently, they contribute to answering the second part of the question by first discussing possible means to tackle the challenges offered by legacy code verification and then suggesting a strategy for deductive PhV, together with possible improvements to existing deductive verification methodologies and tools like KeY [2, 7] and VCC [20].

Cleophas et al. [18] discuss how CbC-based development may lead to a deep comprehension of algorithm families, based on the fact that organising the refinements obtained during CbC-based design in a taxonomy leads to a classification of common and varying properties within a family of algorithms and thus to insight in the relations among its elements. They also argue that using taxonomies in the implementation of toolkits, i.e. a library of all variants, for TABASCO [19] has the additional benefit of providing a meaningful starting point for extractive and proactive SPLE. For both, a concrete methodology is presented.

Kleijn et al. [9] consider systems of systems, represented as team automata, whose components interact by the synchronised execution of common actions [8]. They study conditions for the compatibility of components, defined as being free from message loss and deadlocks, relative to notions of synchronisation other than mandatory synchronised execution. They focus on various kinds of master-slave synchronisations, which require input actions (for ‘slaves’) to be driven by output actions (from ‘masters’). Team automata composed according to this notion of synchronisation are exemplified and studied in some detail, including an extensive discussion of (potential) applications.

Legay et al. [34] introduce DynBLTL, an extension of time-bounded LTL [16], as a new 3-valued logic specifically aimed to reason over dynamically evolving software architectures. Here dynamism is understood as allowing components to be removed, added or (re)connected differently. The third value (undefined) is used to deal with components that are absent in a system configuration being evaluated. The semantics is that of (un)timed traces of graphs seen as snapshots of the architecture at a specific moment of computation. One can quantify over all connections and components of a specific type or a specific component or connection, and the modalities allow to reasoning over a bounded number of steps or a bounded amount of time. Since the number of components is unknown upfront, SMC by means of an integration into the PLASMA statistical model checker [30] is an obvious choice. An example illustrates the approach.

Méry et al. [14] aim to show how CbC and PhV can possibly be combined in a productive way by describing a general framework that integrates the following two different approaches to software verification: program refinement as supported by Event-B [1] and program verification as supported by the Spec# programming system [5]. In particular, they describe a plug-in for Event-B’s RODIN toolset that is able to automatically transform a given abstract Event-B specification into a recursive algorithm that is correct-by-construction and which can be directly translated into executable code.

Schaefer et al. [28] investigate the feasibility of generalizing the concept of proof-carrying code to proof-carrying apps as a means to verify extensible software platforms at deployment time. Rather than global safety policies, contracts are used to specify functional properties of the API of the base software platform, leaving it to the provider of the extension to verify that all API calls adhere to the contract. The resulting proof artefacts are used at deployment time to allow proof checking. After discussing the criteria that enable a verification technique for contract-based deployment-time verification, the applicability of deductive verification with KeY [2, 7] and data-flow analyses with Soot [31] to the proof-carrying apps scenario is examined for a simple Java implementation.

De Vink et al. [10] illustrate the idea of supervisory controller synthesis for SPLE by applying the CIF 3 toolset [11] to an example. They show how to automatically synthesise an SPL model (in the form of an automaton for each valid product of the SPL) starting from a so-called attributed feature model, component behaviour models associated with the features, and additional behavioural requirements (like state invariants, event orderings, and guards on events). The resulting CIF 3 model then satisfies all feature-related constraints as well as all behavioural requirements, by construction. Further behavioural properties can be verified by exporting such SPL models in the input format of the mCRL2 model checker [21].

Beyer [13] outlines a few existing verification approaches that, in an attempt to try to increase the impact of formal verification, combine the advantages of automatic verification techniques and interactive verification techniques. The former, which usually expect the user to set the parameters while the prover computes the necessary invariants and the proof, work well for large systems, whereas the latter, which usually expect the user to provide invariants while the prover establishes a formal correctness proof, work well for sophisticated specifications.

## References

1. J.-R. Abrial. *Modeling in Event-B: System and Software Engineering*. Cambridge University Press, 2010.
2. W. Ahrendt *et al.* The KeY Platform for Verification and Analysis of Java Programs. In D. Giannakopoulou and D. Kroening, editors, *Proceedings of the 6th International Conference on Verified Software: Theories, Tools and Experiments (VSTTE'14)*, volume 8471 of *LNCS*, pages 55–71. Springer, 2014.
3. C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
4. T. Ball and S. K. Rajamani. The SLAM Project: Debugging System Software via Static Analysis. *ACM SIGPLAN Notices*, 37(1):1–3, Jan 2002.
5. M. Barnett, K. R. M. Leino, and W. Schulte. The Spec# Programming System: An Overview. In G. Barthe, L. Burdy, M. Huisman, J.-L. Lanet, and T. Muntean, editors, *Proceedings of the International Conference on Construction and Analysis of Safe, Secure, and Interoperable Smart Devices (CASSIS'04)*, volume 3362 of *LNCS*, pages 49–69. Springer, 2005.

6. B. Beckert, T. Bormer, and D. Grahl. Deductive Verification of Legacy Code. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
7. B. Beckert, R. Hähnle, and P. H. Schmitt. *Verification of Object-Oriented Software: The KeY Approach*. Springer, 2007.
8. M. H. ter Beek and J. Kleijn. Team Automata Satisfying Compositionality. In K. Araki, S. Gnesi, and D. Mandrioli, editors, *Proceedings of the 12th International Symposium on Formal Methods (FM'03)*, volume 2805 of LNCS, pages 381–400. Springer, 2003.
9. M. H. ter Beek, J. Kleijn, and J. Carmona. Conditions for Compatibility of Components: The case of masters and slaves. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
10. M. H. ter Beek, M. A. Reniers, and E. P. de Vink. Supervisory Controller Synthesis for Product Lines using CIF 3. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
11. D. A. van Beek, W. J. Fokkink, D. Hendriks, A. Hofkamp, J. Markovski, J. M. van de Mortel-Fronczak, and M. A. Reniers. CIF 3: Model-Based Engineering of Supervisory Controllers. In E. Ábrahám and K. Havelund, editors, *TACAS'14*, volume 8413 of LNCS, pages 575–580. Springer, 2014.
12. A. Bertolino, P. Inverardi, and H. Muccini. Software architecture-based analysis and testing: a look into achievements and future challenges. *Computing*, 95(8):633–648, 2013.
13. D. Beyer. Partial Verification and Intermediate Results as a Solution to Combine Automatic and Interactive Verification Techniques. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
14. Z. Cheng, D. Méry, and R. Monahan. On two Friends for getting Correct Programs: Automatically Translating Event-B Specifications to Recursive Algorithms in Rodin. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
15. E. M. Clarke, E. A. Emerson, and J. Sifakis. Model checking: algorithmic verification and debugging. *Communications of the ACM*, 52(11):74–84, 2009.
16. E. M. Clarke, J. R. Faeder, C. J. Langmead, L. A. Harris, S. K. Jha, and A. Legay. Statistical Model Checking in BioLab: Applications to the Automated Analysis of T-Cell Receptor Signaling Pathway. In M. Heiner and A. M. Uhrmacher, editors, *Proceedings of the 6th International Conference on Computational Methods in Systems Biology (CMSB'08)*, volume 5307 of LNCS, pages 231–250. Springer, 2008.
17. E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. MIT Press, 1999.
18. L. Cleophas, D. G. Kourie, V. Pieterse, I. Schaefer, and B. W. Watson. Correctness-by-Construction  $\wedge$  Taxonomies  $\Rightarrow$  Deep Comprehension of Algorithm Families. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
19. L. G. Cleophas, B. W. Watson, D. G. Kourie, A. Boake, and S. A. Obiedkov. TABASCO: using concept-based taxonomies in domain engineering. *South African Computer Journal*, 37:30–40, 2006.

20. E. Cohen *et al.* VCC: A Practical System for Verifying Concurrent C. In S. Berghofer, T. Nipkow, C. Urban, and M. Wenzel, editors, *Proceedings of the 22nd International Conference on Theorem Proving in Higher Order Logic (TPHOLs'09)*, volume 5674 of *LNCS*, pages 23–42. Springer, 2009.
21. S. Cranen, J. F. Groote, J. J. A. Keiren, F. P. M. Stappers, E. P. de Vink, W. Weselink, and T. A. C. Willemse. An Overview of the mCRL2 Toolset and Its Recent Advances. In N. Piterman and S. A. Smolka, editors, *TACAS'13*, volume 7795 of *LNCS*, pages 199–213. Springer, 2013.
22. E. W. Dijkstra. A constructive approach to the problem of program correctness. *BIT Numerical Mathematics*, 8(3):174–186, 1968.
23. J.-C. Filliâtre. Deductive software verification. *International Journal on Software Tools for Technology Transfer*, 13(5):397–403, 2011.
24. A. Hall. Correctness by Construction: Integrating Formality into a Commercial Development Process. In L. Eriksson and P. A. Lindsay, editors, *FME 2002*, volume 2391 of *LNCS*, pages 224–233. Springer, 2002.
25. A. Hall and R. Chapman. Correctness by Construction: Developing a Commercial Secure System. *IEEE Software*, 19(1):18–25, Jan/Feb 2002.
26. S. Hangal and M. S. Lam. Tracking Down Software Bugs Using Automatic Anomaly Detection. In *Proceedings of the 24th International Conference on Software Engineering (ICSE'02)*, pages 291–301. ACM, 2002.
27. C. A. R. Hoare. Proof of a Program: FIND. *Communications of the ACM*, 14(1):39–45, 1971.
28. S. Holthusen, M. Nieke, T. Thüm, and I. Schaefer. Proof-Carrying Apps: Contract-Based Deployment-Time Verification. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
29. D. Hovemeyer and W. Pugh. Finding Bugs is Easy. *ACM SIGPLAN Notices*, 39(12):92–106, Dec 2004.
30. C. Jegourel, A. Legay, and S. Sedwards. A Platform for High Performance Statistical Model Checking – PLASMA. In C. Flanagan and B. König, editors, *Proceedings of the 18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'12)*, volume 7214 of *LNCS*, pages 498–503. Springer, 2012.
31. P. Lam, E. Bodden, O. Lhoták, and L. Hendren. The Soot framework for Java program analysis: a retrospective. In *Cetus Users and Compiler Infrastructure Workshop (CETUS'11)*, 2011.
32. A. P. Mathur. *Foundations of Software Testing*. Addison-Wesley, 2nd edition, 2014.
33. M. Pezzè and M. Young. *Software Testing and Analysis: Process, Principles and Techniques*. Wiley, 2007.
34. J. Quilbeuf, E. Cavalcante, L.-M. Traonouez, F. Oquendo, T. Batista, and A. Legay. A Logic for the Statistical Model Checking of Dynamic Software Architectures. In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.
35. J. A. Robinson and A. Voronkov, editors. *Handbook of Automated Reasoning*. MIT Press, 2001.
36. B. W. Watson, D. G. Kourie, I. Schaefer, and L. G. Cleophas. Correctness-by-Construction and Post-hoc Verification: A Marriage of Convenience? In T. Margaria and B. Steffen, editors, *Proceedings of the 6th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA'16)*, LNCS. Springer, 2016.