

# The JKIND Model Checker

Andrew Gacek<sup>1</sup>, John Backes<sup>1</sup>, Mike Whalen<sup>2</sup>, Lucas Wagner<sup>1</sup>, and Elaheh Ghassabani<sup>2</sup>

<sup>1</sup> Rockwell Collins

andrew.gacek@gmail.com, john.backes@gmail.com,  
lucas.wagner@rockwellcollins.com

<sup>2</sup> University of Minnesota

ghass013@umn.edu, mwwhalen@umn.edu



**Abstract.** JKIND is an open-source industrial model checker developed by Rockwell Collins and the University of Minnesota. JKIND uses multiple parallel engines to prove or falsify safety properties of infinite state models. It is portable, easy to install, performance competitive with other state-of-the-art model checkers, and has features designed to improve the results presented to users: *inductive validity cores* for proofs and *counterexample smoothing* for test-case generation. It serves as the back-end for various industrial applications.

## 1 Introduction

JKIND is an open-source<sup>3</sup> industrial infinite-state inductive model checker for safety properties. Models and properties in JKIND are specified in LUSTRE [17], a synchronous data-flow language, using the theories of linear real and integer arithmetic. JKIND uses SMT-solvers to prove and falsify multiple properties in parallel. A distinguishing characteristic of JKIND is its focus on the usability of results. For a proven property, JKIND provides traceability between the property and individual model elements. For a falsified property, JKIND provides options for simplifying the counterexample in order to highlight the root cause of the failure. In industrial applications, we have found these additional usability aspects to be at least as important as the primary results. Another important characteristic of JKIND is that it is designed to be integrated directly into user-facing applications. Written in Java, JKIND runs on all major platforms and is easily compiled into other Java applications. JKIND bundles the Java-based SMTINTERPOL solver and has no external dependencies. However, it can optionally call Z3, YICES 1, YICES 2, CVC4, and MATHSAT if they are available.

## 2 Functionality and Main Features

JKIND is structured as several parallel engines that coordinate to prove properties, mimicking the design of PKIND and KIND 2 [8, 21]. Some engines are

<sup>3</sup> <https://github.com/agacek/jkind>

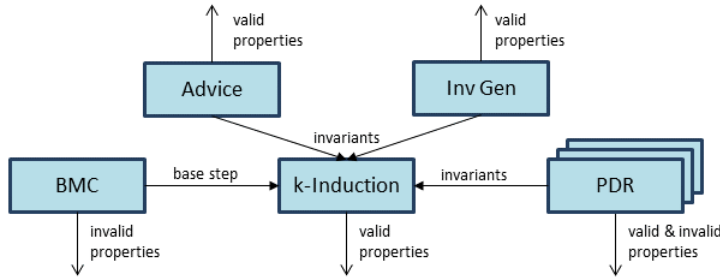


Fig. 1. JKIND engine architecture

directly responsible for proving properties, others aid that effort by generating invariants, and still others are reserved for post-processing of proof or counterexample results. Each engine can be enabled or disabled separately based on the user’s needs. The architecture of JKIND allows any engine to broadcast information to the other engines (for example, lemmas, proofs, counterexamples) allowing straightforward integration of new functionality.

The solving engines in JKIND are shown in Figure 1. The **Bounded Model Checking (BMC)** engine performs a standard iterative unrolling of the transition relation to find counterexamples and to serve as the base case of  $k$ -induction. The BMC engine guarantees that any counterexample it finds is minimal in length. The  **$k$ -Induction** engine performs the inductive step of  $k$ -induction, possibly using invariants generated by other engines. The **Invariant Generation** engine uses a template-based invariant generation technique [22] using its own  $k$ -induction loop. The **Property Directed Reachability (PDR)** engine performs property directed reachability [11] using the implicit abstraction technique [9]. Unlike BMC and  $k$ -induction, each property is handled separately by a different PDR sub-engine. Finally, the **Advice** engine produces invariants based on previous runs of JKIND as described in the next section.

Invariant sharing between the solvers (shown in Figure 1) is an important part of the architecture. In our internal benchmarking, we have found that implicit abstraction PDR performs best when operating over a single property at a time and without use of lemmas generated by other approaches. On the other hand, the invariants generated by PDR and template lemma generation often allow  $k$ -induction, which operates on all properties in parallel, to substantially reduce the verification time required for models with large numbers of properties.

## 2.1 Post Processing and Re-verification

A significant part of the research and development effort for JKIND has focused on post-processing results for presentation and repeated verification of models under development.

**Inductive Validity Cores (IVC).** For a proven property, an inductive validity core is a subset of LUSTRE equations from the input model for which the property still holds [13, 14]. Inductive validity cores can be used for traceability

from property to model elements and determining coverage of the model by a set of properties [15]. This facility can be used to automatically generate traceability and adequacy information (such as traceability matrices [12] important to the certification of safety-critical avionics systems [26]). The IVC engine uses a heuristic algorithm to efficiently produce minimal or nearly minimal cores. In a recent experiment over a superset of the benchmark models described in the experiment in Section 3, we found that our heuristic IVC computation added 31% overhead to model checking time, and yielded cores approximately 8% larger than the guaranteed minimal core computed by a very expensive “brute force” algorithm. As a side-effect, the IVC algorithm also minimizes the set of invariants used to prove a property and emits this reduced set to other engines (notably the *Advice* engine, described below).

**Smoothing.** To aid in counterexample understanding and in creating structural coverage tests that can be more easily explained, JKIND provides an optional post-processing step to minimize the number of changes to input variables—*smoothing* the counterexample. For example, applied to 129 test cases generated for a production avionics flight control state machine, smoothing increased runtime by 40% and removed 4 unnecessary input changes per test case on average. The smoothing engine uses a MAXSAT query over the original BMC-style unrolling of the transition relation combined with weighted assertions that each input variable does not change on each step. The MAXSAT query tries to satisfy all of these weighted assertions, but will break them if needed. This has the effect of trying to hold all inputs constant while still falsifying the original property and only allowing inputs to change when needed. This engine is only available with SMT-solvers that support MAXSAT such as YICES 1 and Z3.

**Advice.** The advice engine saves and re-uses the invariants that were used by JKIND to prove the properties of a model. Prior to analysis, JKIND performs model slicing and flattening to generate a flat transition-relation model. Internally, invariants are stored as a set of proven formulas (in the LUSTRE syntax) over the variables in the flattened model. An *advice* file is simply the emitted set of these invariant formulas. When a model is loaded, the formulas are loaded into memory. Formulas that are no longer syntactically or type correct are discarded, and the remaining set of formulas are submitted as an initial set of possible invariants to be proved via  $k$ -induction. If they are proved, they are passed along to other engines; if falsified, they are discarded. Names constructed between multiple runs of JKIND are stable, so if a model is unchanged, it can be usually be re-proved quickly using the invariants and  $k$ -induction. If the model is slightly changed, it is often the case that most of the invariants can be re-proved, leading to reduced verification times.

If the IVC engine is also enabled, then advice emits a (close to) minimal set of lemmas used for proof; this often leads to faster re-verification (but more expensive initial verification), and can be useful for examining which of the generated lemmas are useful for proofs.




Fig. 2. Performance benchmarks

### 3 Experimental Evaluation

We evaluated the performance of JKIND against KIND 2 [8], ZUSTRE [20], Generalized PDR in Z3 [19], and IC3 in NUXMV [9]. We used the default options for each tool (using `check_invar_ic3` for NUXMV). Our benchmark suite comes from [9] and contains 688 models over the theory of linear integer arithmetic<sup>4</sup>. All experiments were performed on a 64-bit Ubuntu 17.10 Linux machine with a 12-core Intel Xeon CPU E5-1650 v3 @ 3.50GHz, with 32GB of RAM and a time limit of 60 seconds per model.

Performance comparisons are shown in Figure 2. The key describes the number of benchmarks solved for each tool, and the graph shows the aggregate time required for solving, ordered by time required per-problem, ordered independently for each tool. JKIND was able to verify or falsify the most properties, although Z3 was often the fastest tool. Many of the benchmarks in this set are quickly evaluated: Z3 solves the first 400 benchmarks in just over 12 seconds. Due to JKIND’s use of Java, the JVM/JKIND startup time for an empty model is approximately 0.35s, which leads to poor performance on small models<sup>5</sup>. As always, such benchmarks should be taken with a large grain of salt. In [8], a different set of benchmarks slightly favored KIND 2, and in [9], NUXMV was the most capable tool. We believe that all the solvers are relatively competitive.

### 4 Integration & Applications

JKIND is the back-end for a variety of user-facing applications. In this section, we briefly highlight a few of these applications and how they employ the features discussed previously.

<sup>4</sup> <https://es.fbk.eu/people/griggio/papers/tacas14-ic3ia.tar.bz2>. Note that we removed 263 duplicate benchmarks from the original set.

<sup>5</sup> Without startup time, the curve for JKIND is close to the curve for ZUSTRE

(1) The *Specification and Analysis of Requirements* (SPEAR) tool is an open-source tool for prototyping and analysis of requirements [12]. Starting from a set of formalized requirements, SPEAR uses JKIND to determine whether or not the requirements meet certain *properties*. It uses IVCs to create a traceability matrix between requirements and properties, highlighting unused requirements, over-constrained properties, and other common problems. SPEAR also uses JKIND with smoothing for test-case generation using the Unique First Cause criteria [28].

(2) The *Assume Guarantee Reasoning Environment* (AGREE) tool is an open-source compositional verification tool that proves properties of hierarchically-composed models in the Architectural Analysis and Design Language (AADL) language [3, 10, 23]. AGREE makes use of multiple JKIND features including smoothing to present clear counterexamples, IVC to show requirements traceability, and counterexample generation to check the consistency of an AADL component’s contract. AGREE also uses JKIND for test-case generation from component contracts.

(3) The *Static IMPerative AnaLyzer* (SIMPAL) tool is an open-source tool for compositional reasoning over software [27]. SIMPAL is based on LIMP, a LUSTRE-like imperative language with extensions for control flow elements, global variables, and a syntax for specifying preconditions, postconditions, and global variable interactions of preexisting components. SIMPAL translates LIMP programs to an equivalent LUSTRE representation which is passed to JKIND to perform assume-guarantee reasoning, reachability, and viability analyses.

(4) JKIND is also used by two proprietary tools used by product areas within Rockwell Collins. The first is a *Mode Transition Table* verification tool used for the complex state machines which manage flight modes of an aircraft. JKIND is used to check properties and generate tests for mode and transition coverage from LUSTRE models generated from the state machines. IVCs are used to establish traceability, i.e. which transitions are covered by which properties. The second is a *Crew Alerting System* MC/DC test-case generation tool for a proprietary domain-specific language used for messages and alerts to airplane pilots. Smoothing is very important in this context as test cases need to be run on the actual hardware where timing is not precisely controllable. Thus, test cases with a minimum of changes to the inputs are ideal.

## 5 Related Work

JKIND is one of a number of similar infinite-state inductive model checkers including KIND 2 [8], NUXMV [9], Z3 with generalized PDR [19], and ZUSTRE [20]. They operate over a transition relation described either as a LUSTRE program (KIND 2, JKIND, and ZUSTRE), an extension of the SMV language (NUXMV), or as a set of Horn clauses (Z3). Each tool uses a portfolio-based solver approach, with NUXMV, JKIND, and KIND 2 all supporting both  $k$ -induction and a variant of PDR/IC3. NUXMV also supports guided reachability and  $k$ -liveness. Other tools such as ESBMC-DEPTHK [25], VVT [4]

CPACHECKER, [5], CPROVER [7] use similar techniques for reasoning about C programs.

We believe that the JKIND IVC support is similar to *proof-core* support provided by commercial hardware model checkers: Cadence Jasper Gold and Synopsys VC Formal [1,2,18]. The proof-core provided by these tools is used for internal coverage analysis measurements performed by the tools. Unfortunately, the algorithms used in the commercial tool support are undocumented and performance comparisons are prohibited by the tool licenses, so it is not possible to compare performance on this aspect.

Previous work has been done on improving the quality of counterexamples along various dimensions similar to the JKIND notion of *smoothing*, e.g. [16,24]. Our work is distinguished by its focus on minimizing the number of deltas in the input values. This metric has been driven by user needs and by our own experiences with test-case generation.

There are several tools that support reuse or exchange of verification results, similar to our *advice* feature. Recently, there has been progress on standardized formats [6] of exchange between analysis tools. Our current advice format is optimized for use and performance with our particular tool and designed for re-verification rather than exchange of partial verification information. However, supporting a standardized format for exchanging verification information would be a useful feature for future use.

## 6 Conclusion

JKIND is similar to a number of other solvers that each solve infinite state sequential analysis problems. Nevertheless, it has some important features that distinguish it. First, a focus on quality of feedback to users for both valid properties (using IVCs) and invalid properties (using smoothing). Second, it is supported across all major platforms and is straightforward to port due to its implementation in Java. Third, it is small, modular, and well-architected, allowing straightforward extension with new engines. Fourth, it is open-source with a liberal distribution license (BSD), so it can be adapted for various purposes, as demonstrated by the number of tools that have incorporated it.

### Acknowledgments

The work presented here was sponsored by DARPA as part of the HACMS program under contract FA8750-12-9-0179.

## References

1. Cadence JasperGold Formal Verification Platform. <https://www.cadence.com/>
2. Synopsys VC Formal Platform. <https://www.synopsys.com>
3. Backes, J., et al.: Requirements analysis of a quad-redundant flight control system. In: NFM (2015)
4. Beyer, D., Dangl, M.: SMT-based software model checking: An experimental comparison of four algorithms. In: VSTTE (2016)

5. Beyer, D., et al.: Boosting k-induction with continuously-refined invariants. In: CAV (2015)
6. Beyer, D., et al.: Correctness witnesses: Exchanging verification results between verifiers. In: FSE. pp. 326–337 (2016)
7. Brain, M., et al.: Safety verification and refutation by k-invariants and k-induction. In: SAS (2015)
8. Champion, A., et al.: The Kind 2 model checker. In: CAV (2016)
9. Cimatti, A., et al.: IC3 modulo theories via implicit predicate abstraction. In: TACAS (2014)
10. Cofer, D., et al.: Compositional verification of architectural models. In: NFM (2012)
11. Eén, N., et al.: Efficient implementation of property directed reachability. In: FM-CAD (2011)
12. Fifarek, A.W., et al.: SpeAR v2.0: Formalized past LTL specification and analysis of requirements. In: NFM (2017)
13. Ghassabani, E., et al.: Efficient generation of all minimal inductive validity cores. In: FMCAD (2017)
14. Ghassabani, E., et al.: Efficient generation of inductive validity cores for safety properties. In: FSE (2016)
15. Ghassabani, E., et al.: Proof-based coverage metrics for formal verification. In: ASE (2017)
16. Groce, A., Kroening, D.: Making the most of BMC counterexamples. In: Proceedings of BMC 2004. ENTCS, vol. 119, pp. 67–81. Elsevier (2005)
17. Halbwachs, N., et al.: The synchronous dataflow programming language Lustre. IEEE (1991)
18. Hanna, Z., et al.: Formal verification coverage metrics for circuit design properties (2015), <https://www.google.com/patents/US20150135150>
19. Hoder, K., Bjørner, N.: Generalized property directed reachability. In: SAT 2012. pp. 157–171. Berlin, Heidelberg (2012)
20. Kahsai, T., Bourbouh, H.: Zustre product home page. <https://github.com/coco-team/zustre>
21. Kahsai, T., Tinelli, C.: Pkind: A parallel k-induction based model checker. In: PDMC (2011)
22. Kahsai, T., et al.: Incremental verification with mode variable invariants in state machines. In: NFM (2012)
23. Murugesan, A., et al.: Compositional verification of a medical device system. In: HILT (2013)
24. Ravi, K., Somenzi, F.: Minimal assignments for bounded model checking. In: Jensen, K., Podelski, A. (eds.) Tools and Algorithms for the Construction and Analysis of Systems. pp. 31–45. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
25. Rocha, H., et al.: Model checking embedded C software using k-induction and invariants. In: SBESC (2015)
26. RTCA DO-178C: Software Considerations in Airborne Systems and Equipment Certification (2011)
27. Wagner, L., et al.: SIMPAL: A compositional reasoning framework for imperative programs. In: SPIN (2017)
28. Whalen, M.W., et al.: Coverage metrics for requirements-based testing. In: ISSSTA (2006)