

A Program Certification Assistant Based on Fully Automated Theorem Provers

Ewen Denney Bernd Fischer

USRA/RIACS, NASA Ames Research Center
Moffett Field, CA 94035, USA

{edenney,fisch}@email.arc.nasa.gov

Abstract

We describe a *certification assistant* to support formal safety proofs for programs. It is based on a graphical user interface that hides the low-level details of first-order automated theorem provers while supporting limited interactivity: it allows users to customize and control the proof process on a high level, manages the auxiliary artifacts produced during this process, and provides traceability between the proof obligations and the relevant parts of the program. The certification assistant is part of a larger program synthesis system and is intended to support the deployment of automatically generated code in safety-critical applications.

1 Introduction

Program verification remains one of the most promising applications of theorem proving, and both fully automatic and interactive provers have been used in verification projects. However, program verification has not lived up to its early promises and is not yet applied routinely in software development. This has a variety of reasons, ranging from the technical difficulties the task still poses for theorem provers, to problems in designing appropriate interfaces, and the cultural changes that are necessary in the software development process itself.

In this paper, we describe a user interface we have developed for the application of fully automatic first-order theorem provers (ATPs) in formal *program certification*, a limited and thus more tractable variant of full program verification. It uses the same basic technology but is concerned only with safety-relevant aspects of a software system rather than the complete system behavior. This limitation enables the successful application of ATPs. Program certification is based on the idea that mathematical proofs of the individual safety properties can be regarded as certificates which can be subjected to external scrutiny and related to the relevant safety-critical parts of the software system. This particular purpose of the proofs requires a dedicated application-oriented interface rather than a proof-oriented interface—in effect, we need a *certification assistant* rather than a proof assistant.

The role as certification assistant puts the user interface under the influence of two competing design principles. On the one hand, it has to “hide” the theorem provers from unsuspecting software engineers. On the other, it has to be open in order to provide coarse-grained control of the certification process, to maintain traceability between the different artifacts (in particular, source code and verification conditions), and to ensure trust in the entire certification process. As a consequence of this duality, the interface cannot completely separate the ATP from the rest of the certification environment. In particular, it needs to support the interpretation of both the input to and the output from the ATP in terms of the application.

Our work on certification emerged from an ongoing project on *automated program-synthesis*. We have developed two synthesis systems for the domains of scientific data analysis [FS03] and state estimation [WS05], which can generate code for safety-critical application areas like spacecraft guidance, navigation, and control. Process standards such as DO-178B [RTC92], however, require that all safety-critical software be certified to a high degree of confidence. Our goal is thus to integrate our synthesis tools with a dedicated certification environment so that end-users can understand and trust the generated code more easily. We adopt a browser paradigm so that users can inspect the code and interact with the underlying prover, while being shielded from the low-level minutiae. Although this is still work in progress, we believe that it offers potential for increasing acceptance of code generators in safety-critical domains at NASA.

In Section 2, we introduce our automated synthesis and verification systems and describe the certification problem we address. The system architecture has a direct bearing on the certification interface, which is described in Section 3. Although the system is fully automatic, users have the option of controlling the proof process by selecting and parameterizing different provers, and inspecting the logs of prover sessions, including the proofs themselves. This is described in Sections 3.1 and 3.2, respectively. In Section 3.3, we describe the verification condition browser, which is used to relate proof obligations to the synthesized code. Section 4 describes related work on interfaces for prover-based verification, and Section 5 outlines our future plans.

2 Background and System Architecture

Figure 1 shows the overall architecture of our extended program synthesis system, which comprises three classes of components: the original synthesis system, the certification extensions, and corresponding document generation extensions. We will describe these components in some more detail in the following.

2.1 Program Synthesis

Traditionally, program synthesis has followed the proofs-as-programs paradigm: the problem is specified as a conjecture in a suitable logic, an interactive theorem prover like Coq, Isabelle, or NuPRL is used to construct a proof, and a func-




Figure 1: Certifiable program synthesis: System architecture

tional program is extracted from that proof and then translated into the target environment. However, this traditional, purely deductive approach to program synthesis is notoriously difficult to scale up to large problems (cf. [AB01]) and full automation has remained elusive. We thus follow a *schema-based* synthesis approach that combines deductive reasoning with techniques from generative programming. Most of the components described here are hidden inside the synthesis system box in Figure 1.

Problem Specifications. Since schema-based synthesis does not require a logical conjecture as starting point for a proof, the code derivation can begin with a specification in a more application-oriented *domain-specific language*. The details of the language obviously depend on the domain of the synthesis system, but in general it combines some target language constructs (e.g., declarations) with established scientific and engineering notations (e.g., differential equations). This allows a concise and fully declarative formulation of the problem together with some details of the desired configuration and architecture of the code to be generated.

Schemas. A *schema* is a parameterized code fragment (i.e., template) together with a set of constraints that determine whether the schema is applicable and how the parameters can be instantiated. The constraints are formulated as conditions on a problem model, which allows the problem structure to directly guide the application of the schemas and thus constrain the search space. The parameters are instantiated by the synthesis engine, either directly on schema application or by recursive calls with a modified problem. The schemas are organized hierarchically into a *schema library* which further constrains the search space. Schemas represent both fundamental building blocks (i.e., algorithms) and solution methods (i.e., theorems) of the domain; they are thus similar to the lemmas used in interactive systems but they can contain explicit calls to a meta-programming kernel in order to construct the code fragments.

Symbolic Computations. Symbolic computations are used to support schema instantiation and code optimization. The core of the symbolic subsystem is a small rewrite engine which supports associative-commutative operators and explicit contexts. It thus allows contextual rules as for example $x/x \rightarrow_{C \vdash x \neq 0} 1$ where $\rightarrow_{C \vdash x \neq 0}$ means "rewrites to, provided $x \neq 0$ can be proven from the current context C ." Expression simplification and symbolic differentiation, similar to those in Mathematica, are implemented on top of the rewrite engine. The basic rules are straightforward; however, vectors and matrices require careful formalizations, and some rules also require explicit meta-programming, e.g., when bound variables are involved.

Intermediate Code. The code fragments in the schemas are formulated in an imperative intermediate language. This is essentially a "sanitized" variant of C (i.e., no pointers, side effects in expressions etc.); however, it also contains a number of domain-specific constructs like vector/matrix operations, finite sums, and convergence loops.

Optimization. Straightforward schema instantiation and composition produces suboptimal code; worse, many of the suboptimalities cannot be removed completely using a separate, after-the-fact optimization phase. Schemas can thus explicitly trigger large-scale optimizations which take into account information from the synthesis process. For example, all numeric routines restructure the goal expression using code motion, common sub-expression elimination, and memoization; since the schemas know the goal variables, no dataflow analysis is required to identify invariant sub-expressions, and code can be moved around aggressively, even across procedure borders.

Code Generation. In a final step, the optimized intermediate code is translated into code tailored for a specific run-time environment. We currently have code generators for the Octave and Matlab environments, and can also produce standalone Ada, C, and Modula-2 code. Each code generator employs one rewrite system to eliminate the constructs of the intermediate language which are not supported by the target environment ("desugaring") and a second rewrite system to clean up the desugared code; most rules are shared between the different code generators.

AUTOBAYES and AUTOFILTER. So far we have built two domain-specific synthesis systems following the schema-based approach outlined above. AUTOBAYES [FS03] works in the scientific data analysis domain and generates parameter learning programs, while AUTOFILTER [WS05] generates state estimation code based on variants of the Kalman filter algorithm. Both systems share a large common core (e.g., symbolic subsystem, certification subsystem, and code generators) but have their individual schema libraries. They are implemented in SWI-Prolog and together comprise approximately 100 kloc. Both systems work fully automatically and can generate code of considerable size and complexity (approximately 1500 loc with deeply nested loops) within a few seconds.

2.2 Certification

Unlike purely deductive approaches, schema-based synthesis cannot ensure “correctness-by-construction”. Since formally verifying the entire system is unfeasible, we instead validate each generated program individually; furthermore, we concentrate on specific aspects of program safety (e.g., memory safety). The core idea is that the schemas can be extended to simultaneously generate code *and all required annotations* such that a verification condition generator can produce proof obligations which are then discharged using an automated theorem prover. The proofs, which can be validated by an automated proof checker or prepared for human inspection, then serve as *certificates*. This approach is tractable because the synthesis system has full knowledge about the form the generated code will take and the specific safety aspect that is to be certified, so that it can generate the appropriate annotations. However, our certification approach is not necessarily tied to synthesis and the annotations could in principle also be added manually.

Safety Policies. A *safety policy* is a set of Hoare-style proof rules and auxiliary definitions which are designed to show that “a program does not go wrong,” i.e., satisfies the safety property of interest [WSF02, DF03]. Safety policies can be used to enforce both *language-specific* properties which can be expressed in terms of the constructs of the underlying programming language, itself, and are thus sensible for any program in the language, as well as *domain-specific* properties, which typically relate to high-level concepts outside the language (e.g., matrix multiplication).

We currently support five different safety policies. Array-bounds safety requires each access to an array element to be within the declared bounds of the array. Variable initialization-before-use ensures that each variable or individual array element has been assigned a defined value before it is used. Both are typical examples of language-specific properties. For the data analysis domain, we can guarantee vector-norm safety (i.e., probability vectors add up to one), and for the state estimation domain we can check proper sensor input usage (i.e., all input variables are used in the computation of the filter output) and matrix symmetry (i.e., covariance matrices are not skewed).

Annotated Code. The annotations are part of the schema and thus are instantiated in parallel with the code fragments; further annotations are introduced by the desugaring steps of the code generation phase. The annotations contain local information in the form of logical pre- and post-conditions and loop invariants, which is then propagated through the code.

VCG. The fully annotated code is then processed by a weakest precondition verification condition generator (VCG), which applies the Hoare-rules of the safety policy in order to generate verification conditions (VCs). The VCG has been designed to be “correct-by-inspection”, i.e., to be sufficiently simple that it is straightforward to see that it correctly implements the rules of the logic. Hence, it does not implement any optimizations, such as structure sharing on the VCs or even apply any simplifications. As usual, the VCG works backwards through the code and verification conditions are generated at each line that can

potentially violate the safety policy.

Simplification. By design of the VCG, the generated VCs are quite complex; hence, they need to be simplified before they can be discharged by an ATP. The certification extension thus re-uses the rewrite engine of the synthesizer together with a dedicated set of rewrite rules. Details can be found in [DFS04a].

ATP. For our purposes, an ATP is a search procedure which applies the inference rules of its calculus until it either finds a proof or fails because none of the rules are applicable. In order to handle extra-logical operations (as, for example, arithmetic functions), the ATP needs an additional *domain theory* that specifies their intended meaning as axioms. The provers use a set of core axioms, together with a collection of dynamically generated axioms, depending on the particular proof task.

Proof Checking. As an alternative to formally verifying the ATPs, they can be extended to generate sufficiently detailed proofs which can then be independently checked by a small and thus verifiable algorithm. However, due to the lack of a standardized format, (and various other reasons [DFS04b]) there are almost no proof checkers for high-performance ATPs, in contrast to the situation for interactive higher-order provers. We have linked our system to the only exceptions we are aware of: the IVY system [MS00], which is based on Otter, and the GDV verifier [SB05].

Trusted Components. Similarly to proof carrying code [NL98], we distinguish between trusted and untrusted components, shown in Figure 1 in red (dark grey) and blue (light grey), respectively. Components are called *trusted*—and must thus be correct—if any errors in them can compromise the assurance provided by the overall system. *Untrusted* components, on the other hand, are not crucial to the assurance because their results are double-checked by at least one trusted component. In particular, the correctness of the certification system does not depend on the correctness of the two largest components: the synthesizer, and the theorem prover; instead, we need only trust the safety policy, the VCG, the domain theory, and the proof checker.

2.3 Document Generation

The basic idea behind our certification approach can also be extended to human-readable documentation. The schemas contain text templates that are instantiated and composed together with the code fragments; these auto-generated comments explain selected parts of the algorithm, give detailed derivations of mathematical formulas, and relate program constructs and variable names back to the specification.

In addition to the commented code, we can also generate a standardized software design document that contains interface descriptions, administrative information (names of files, versions, etc.), specific input and output constraints, and synthesis and compiler warnings. The document is hyperlinked to the input specification, the code, and any other intermediate artifacts generated during synthesis. Since the design document is generated at synthesis time, it can




Figure 2: Certification assistant: start-up view

include design details which would be difficult, if not impossible, to infer after the code has been generated.

3 The Certification Assistant

The certification assistant serves three main purposes. First, it allows users to customize and control the proof process on a high level. Second, it manages the auxiliary artifacts produced by this process. Third, it provides traceability between the VCs and the relevant parts of the program.

The interface mainly uses straightforward HTML as underlying technology. This is augmented with some scripting code to support the VC-browsing described in Section 3.3. The certification assistant consists of a few supporting shell scripts to control the provers, some boilerplate CGI and HTML code and a number of PHP files that are auto-generated together with the target code; in total, this amounts to approximately 2000 lines. The generation of the PHP files and the customization of the certification assistant itself is triggered by a simple command line option of the synthesis system.

Figure 2 shows the startup view of the certification assistant, after the code and supporting files have been generated. The window is split into three different areas. The left half contains a hyperlinked version of the generated code; the line numbers are used as labels by the VC linker. The right half contains a simple prover control panel and the list of verification conditions below that. Initially,




Figure 3: Certification assistant: updated view with prover results

no information is available about the proof status of any of the verification conditions.

3.1 Prover Control

In contrast to interactive proof systems like Coq, Isabelle, or NuPRL, the user interaction here is only concerned with the parameterization of the proof process, but not with application of individual tactics, and the prover control panel reflects this restricted interaction style. Different drop-down menus allow the user to select the theorem prover and choose between various predefined axiomatizations of the domain theory that are to be used for the proof attempts, and the level of evidence (i.e., proof status, prover logs, or full proofs) that the prover is required to supply. In addition, the user can specify the time limit for an individual task and any prover-specific parameters that will be passed along unchecked. The certification process is then started by selecting any or all of the verification conditions from the list and sending the request, which the certification assistant passes to the selected prover.

Figure 3 shows the view updated with the prover results: For each verification condition, the current proof status is displayed, together with the elapsed proof time, and a link to its location in the source code (cf. Section 3.3). If evidence is produced, the status contains a link to the evidence (cf. Section 3.2). If some proof attempts fail (e.g., the verification condition *quaternion_ds1_init_0034*), the user can resend a request with different settings or

with a different ATP.

The certification assistant serves as common interface to different first-order ATPs in a similar spirit to how Proof General [Asp00] serves as common interface for different interactive higher-order provers. However, due to the limited interaction and the black-box style integration, the protocol requirements to link to a prover are much simpler. In particular, we can use the TPTP syntax [SS03] as common notation which is understood by most targeted ATPs and can be easily translated for the others. Hence, only straightforward control scripts are required for the integration; we are currently also in the process of replacing these by the control scripts used in the annual ATP system competition (CASC) [Sut04]. The certification assistant currently integrates seven different ATPs (DCTP [LS01], E [Sch02], E-Setheo [MI+97], Gandalf [Tam97], Ivy [MS00], Otter [MW97], Spass [Wei03], and Vampire [RV02]), most of them in different versions, that run on the local server. In addition, it also provides a remote link to the SystemOnTPTP proof server at the University of Miami [Sut05], which then acts as a trusted prover component repository.

3.2 Certificate Inspection and Visualization

The certification assistant also provides access to the auxiliary artifacts that are produced during the certification. This includes the intermediate stages in the processing chain (generated axioms, clausal normal form etc.), prover log files, and actual proofs, depending on the required level of evidence. These artifacts can support, or in the absence of a proof collectively serve as, the certificate, and can be inspected as raw text files, or using third-party tools, e.g., the GDV derivation verifier [SB05] and the proof visualizer from the TPTP tool suite [SS03].

3.3 VC Linker and Browser

A VC can fail to be proven for a number of reasons. First, there may of course be an actual safety violation in the code. Second, the (generated) annotations may be insufficient or wrong. Errors can come from any part of the schema, or from the propagation phase: an annotation might not be propagated far enough, or it might be propagated out of scope. Third, the theorem prover may time-out, either due to the size and complexity of the VC, or due to an incomplete domain theory. For certification purposes, however, it is important to distinguish between unsafe programs and any other reasons for failure, and in the case of genuine safety violations, to locate the unsafe parts of the program.

However, manually tracing the VCs back to their source is quite difficult as the verification process is inherently logically complex. The VCs can become very large and go through substantial structural simplifications, after which they are typically [DFS04a] of the form $hyp_1 \wedge \dots \wedge hyp_n \Rightarrow conc$. Here, a hypothesis, hyp , stems either from a loop invariant, an index bound, or a propagated annotation, and the conclusion, $conc$, is either derived from an annotated




Figure 5: Certification assistant: linking from code

which parts of the program are affected, and eventually how the violation can be resolved. This traceability is also mandated by relevant standards such as DO-178B [RTC92].

In practice, safety checks are often carried out during code reviews [NS04], where reviewers look in detail at each line of the code and check the individual safety properties statement by statement. Fortunately, linking works in both directions: clicking on a statement or annotation displays all VCs to which it contributes (i.e., which are labeled with its line number). Figure 5 shows the result of clicking on the label for line 220; the unproven verification condition implies that this line of code has not been cleared yet.

4 Related work

There are a number of program verification systems which use theorem proving, both automated and interactive, as the underlying technology. However, the provers are hidden to varying degrees. Fully automated provers are typically used as black-box components in a tool chain that controls the verification process and determines the form of the interface. Often, the interfaces follow a compiler style, with command line parameters as inputs and error messages as outputs (e.g., ESC/Java [FL⁺02, DNS03]), although an integration with a graphical software development environment is also possible [AB⁺03]. Interactive systems, on the other hand, typically use expressive higher-order logics

to model the entire software within the system, relying on the built-in prover interface to directly control and interpret the verification process.

Perfect Developer [Cro03] is a combined software development and verification environment. The user writes annotated programs in an intermediate language, where the annotations express a correctness specification in the design-by-contract methodology. The programs can then be translated into a variety of target languages, or analyzed using Hoare logic and an automated theorem prover.

Caveat [BC⁺03] offers similar analysis capabilities, although it operates directly on source code (as opposed to an intermediate language), with the advantage that bugs can then be directly located. The tool tries to verify that the code is free of the usual range of safety violations (division by zero, null dereferencing, array out of bounds), based on Hoare logic. If the automated prover fails, the user can start an interactive proof process using Caml as scripting language. The developers' stated aim is to increase the scope for interactivity and control of proof process. Users are also able to manually insert code annotations as required. For debugging purposes, Caveat lets users interpret failed proof obligations, and relate them to their origin. It can also generate counter-examples.

There are very few other approaches to providing explicit traceability for a program synthesis system. Van Baalen and others [vBR⁺98] use origin tracking [vKT93] to indicate how statements in synthesized code relate to the initial problem specification and domain theory. They later built on this to present a documentation generator and XML-based browser interface that generates an explanation for every executable statement in the synthesized program [WB⁺01].

The Proof General [Asp00] generic prover interface aims to shield users from the low-level details of using a theorem prover. It offers a customizable user interface, while adding functionality on top of that provided by the underlying prover. Our aims are similar, although we target a different abstraction level. There are also several developing higher-order proof networks (e.g., Mathweb [FH⁺99]), but it is not yet clear what role they should play in a certification assistant. The SystemOnTPTP proof service [Sut05], which we use, can be seen as a first-order automated equivalent of these.

5 Future Work

Our current efforts focus on extending the certification system in a number of areas. We are currently extending the markup the VCG adds to the proof obligations in two different ways. First, we are replacing the plain line numbers currently used as locations by individual subterm positions, similar to [Fra96]. This allows a more precise tracing. Second, and more significantly, we are extending the VCG to augment the VCs with semantic information concerning the interpretation of the different parts of the formula (e.g., "loop invariant asserted in line X "). This simplifies the interpretation of the VCs for debugging purposes and can also be used to generate high-level descriptions of the VCs.

Since both extensions change the structure of the location information, we now need to correspondingly extend the interface to make use of this information.

We are also extending the interface to include specifications and design documents, thus combining our work on certification with automated safety and design document generation [DV04]. Eventually, we aim to provide full and seamless traceability between specification, design documentation, code, and certificates.

Our broader aim is to develop a certificate management system along the lines of the Programatica project [The03]. This will enable a wide range of additional capabilities, such as support for manual sign-offs on code fragments that violate stated safety policies.

Acknowledgments. Amber Telfer implemented a first version of the VC linking and browsing. Phil Oh helped extend this into the current version of the certification assistant.

References

- [AB01] A. Ayari and D. Basin. “A Higher-order Interpretation of Deductive Tableau”. *JSC*, **31**(5):487–520, May 2001.
- [AB+03] W. Ahrendt, T. Baar, B. Beckert, R. Bubel, M. Giese, R. Hähnle, W. Menzel, W. Mostowski, A. Roth, S. Schlager, and P. H. Schmitt. The KeY Tool. Technical Report TR 2003-05, Göteborg University, 2003.
- [Asp00] D. Aspinall. “Proof General: A Generic Tool for Proof Development”. In *Proceedings of the 6th International Conference on Tools and Algorithms for Construction and Analysis of Systems, Lecture Notes in Computer Science 1785*, pp. 38–42. Springer-Verlag, 2000.
- [BC+03] P. Baudin, G. Canet, A. Pacalet, and D. Schoen. “Caveat: a Tool for Analysis and formal Proof of C Programs”. In *Tool Exhibition Notes, FM 2003: 12th International FME Symposium*, pp. 6–10, Pisa, Italy, 2003.
- [Cro03] D. Crocker. “Perfect Developer: a Tool for Object-Oriented Formal Specification and Refinement”. In *Tool Exhibition Notes, FM 2003: 12th International FME Symposium*, pp. 37–41, Pisa, Italy, 2003.
- [DF03] E. Denney and B. Fischer. “Correctness of Source-Level Safety Policies”. In K. Araki, S. Gnesi, and D. Mandrioli, (eds.), *Proc. FM 2003: Formal Methods, LNCS 2805*, pp. 894–913, Pisa, Italy, September 2003. Springer.
- [DFS04a] E. Denney, B. Fischer, and J. Schumann. “An Empirical Evaluation of Automated Theorem Provers in Software Certification”. In *Proceedings of the IJCAR 2004 Workshop on Empirically Successful First Order Reasoning (ESFOR)*, 2004.

- [DFS04b] E. Denney, B. Fischer, and J. Schumann. "An Empirical Evaluation of Automated Theorem Provers in Software Certification". *International Journal of AI Tools*, 2004. Submitted for publication.
- [DNS03] D. L. Detlefs, G. Nelson, and J. B. Saxe. Simplify: A Theorem Prover for Program Checking. Technical Report HPL-2003-148, HP Labs, 2003.
- [DV04] E. Denney and R. P. Venkatesan. "A Generic Software Safety Document Generator". In *Proceedings of the 10th International Conference on Algebraic Methodology and Software Technology, AMAST'04*, pp. 102–116, Stirling, Scotland, 2004.
- [FH⁺99] A. Franke, S. M. Hess, C. G. Jung, M. Kohlhase, and V. Sorge. "Agent-Oriented Integration of Distributed Mathematical Services". *Journal of Universal Computer Science*, 5(3):156–187, 1999.
-
- [FL⁺02] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. "Extended static checking for Java". In L. J. Hendren, (ed.), *Proc. PLDI 2002*, pp. 234–245, Berlin, Germany, June 17–19 2002. ACM Press. Published as SIGPLAN Notices 37(5).
- [Fra96] R. Fraer. "Tracing the Origins of Verification Conditions". In *Proceedings of the 5th International Conference on Algebraic Methodology and Software Technology, AMAST'96*, pp. 241–255, 1996.
- [FS03] B. Fischer and J. Schumann. "AutoBayes: A System for Generating Data Analysis Programs from Statistical Models". *J. Functional Programming*, 13(3):483–508, May 2003.
-
- [LS01] R. Letz and G. Stenz. "DCTP: A Disconnection Calculus Theorem Prover". In R. Gore, A. Leitsch, and T. Nipkow, (eds.), *Proc. First Intl. Joint Conf. Automated Reasoning, LNAI 2083*, pp. 381–385, Siena, 2001. Springer.
- [MI⁺97] M. Moser, O. Ibens, R. Letz, J. Steinbach, C. Goller, J. Schumann, and K. Mayr. "The Model Elimination Provers SETHEO and E-SETHEO". *JAR*, 18:237–246, 1997.
-
- [MS00] W. McCune and O. Shumsky. "System description: IVY". In D. McAllester, (ed.), *Proc. 17th CADE, LNAI 1831*, pp. 401–405. Springer, 2000.
- [MW97] W. McCune and L. Wos. "Otter—The CADE-13 Competition Incarnations". *JAR*, 18(2):211–220, April 1997.
- [NL98] G. C. Necula and P. Lee. "The Design and Implementation of a Certifying Compiler". In K. D. Cooper, (ed.), *Proc. PLDI 1998*, pp. 333–344, Montreal, Canada, June 17–19 1998. ACM Press. Published as SIGPLAN Notices 33(5).

- [NS04] S. Nelson and J. Schumann. "What makes a Code Review Trustworthy?". In *Proceedings of the Thirty-Seventh Annual Hawaii International Conference on System Sciences (HICSS-37)*. IEEE, 2004.
- [Pro04] The Programatica Project Page, 2004. www.cse.ogi.edu/PacSoft/projects/programatica.
- [RTC92] RTCA Special Committee 167. Software Considerations in Airborne Systems and Equipment Certification. Technical report, RTCA, Inc., December 1992.
- [RV02] A. Riazanov and A. Voronkov. "The Design and Implementation of Vampire". *AI Communications*, 15(2-3):91-110, 2002.
- [SB05] G. Sutcliffe and D. Belfiore. "Semantic Derivation Verification". In I. Russell and Z. Markov, (eds.), *Proceedings of the 18th Florida Artificial Intelligence Research Symposium*. AAAI Press, 2005.
- [Sch02] S. Schulz. "E - A Brainiac Theorem Prover". *Journal of AI Communications*, 15(2/3):111-126, 2002.
- [SS03] G. Sutcliffe and C. Suttner. TPTP Home Page, 2003. www.tptp.org.
- [Sut04] G. Sutcliffe. The CADE-J2 ATP System Competition, 2004. www.tptp.org/CASC/J2/.
- [Sut05] G. Sutcliffe. System on TPTP, 2005. www.tptp.org/cgi-bin/SystemOnTPTPFormMaker.
- [Tam97] T. Tammet. "Gandalf". *JAR*, 18(2):199-204, April 1997.
- [The03] The Programatica Team. "Programatica Tools for Certifiable, Auditable Development of High-assurance Systems in Haskell". In *Proceedings of the High Confidence Software and Systems Conference*, Baltimore, MD, April 2003. Available via [Pro04].
- [vBR⁺98] J. van Baalen, P. Robinson, M. Lowry, and T. Pressburger. "Explaining Synthesized Software". In D. F. Redmiles and B. Nuseibeh, (eds.), *Proc 13th IEEE Conference on Automated Software Engineering*, pp. 240-248, 1998.
- [vKT93] A. van Deursen, P. Klint, and F. Tip. "Origin Tracking". *JSC*, 15(5/6):523-545, 1993.
- [WB⁺01] J. Whittle, J. V. Baalen, J. Schumann, P. Robinson, T. Pressburger, J. Penix, P. Oh, M. Lowry, and G. Brat. "Amphion/NAV: Deductive Synthesis of State Estimation Software". In *Proceedings of the 16th IEEE Conference on Automated Software Engineering*, 2001.
- [Wei03] C. Weidenbach. SPASS Home Page, 2003. <http://spass.mpi-sb.mpg.de>.

[WS05] J. Whittle and J. Schumann. "Automating the Implementation of Kalman Filter Algorithms". *ACM Transactions on Mathematical Software*, 2005. To appear.

[WSF02] M. Whalen, J. Schumann, and B. Fischer. "Synthesizing Certified Code". In L.-H. Eriksson and P. A. Lindsay, (eds.), *Proc. FME 2002: Formal Methods—Getting IT Right, LNCS 2391*, pp. 431–450, Copenhagen, Denmark, July 2002. Springer.