The Importance of Non-theorems and Counterexamples in Program Verification

Graham Steel

School of Informatics,
University of Edinburgh,
Edinburgh, EH8 9LE, Scotland
`graham.steel@ed.ac.uk`
<http://homepages.inf.ed.ac.uk/gsteel>

Abstract. We argue that the detection and refutation of non-theorems, and the discovery of appropriate counterexamples, is of vital importance to the Grand Challenge of a Program Verifier.

1 Introduction

In this essay, we make a case for the inclusion of non-theorem (i.e. incorrect conjecture) detection and counterexample generation as a core theme in the research program of the Grand Challenge of a Program Verifier. We will argue that:

- Research in program verification technology will be hindered if counterexample generation research does not catch up and keep pace. We will give the reasons for this in §2.
- Detecting false conjectures and generating counterexamples to them is a fascinating scientific challenge in itself. We will argue this in §3.
- Deduction based approaches for verification, which offer perhaps the best chance of achieving the goals of the grand challenge, must improve in their handling of non-theorems if they are to compete with model checking approaches, which are already able to give counterexamples to false verification conditions.

Our arguments will be followed in §4 by a look at previous and current research on the topic, including our own efforts in the Mathematical Reasoning Group at the University of Edinburgh.

Since the areas of application are the same, and since the problems of program verification and counterexample generation are in a sense ‘dual’, it seems logical that they should be thought of as part of the same Grand Challenge. Note that we fully support the view that the final goal of the project must be a program verifier, not a system that just finds more and more bugs. However, we will argue in this paper that to achieve this goal, the dual problem of non-theorem detection and counterexample generation must be given due attention.

2 The Importance of Counterexample Generation

Even the most diligent and expert programmer will very rarely write a bug-free first version of a program. The majority of calls to a program verifier will therefore involve an incorrect program. If we want program verification tools to become widely used, they must deal with buggy programs in a competent manner. Simply presenting failed verification conditions or open subgoals is not sufficient, since it leaves the user with no idea whether a bug has been found, or whether the verification system is simply unable to dispose of that particular proof obligation. What is required is a system which can not only detect an incorrect conjecture, but also supply a counterexample to allow the user to locate the flaw.

Non-theorem detection is also important to the internal processes of a verification tool. Automated proof attempts, particularly when induction is involved, will frequently require the conjecturing of lemmas and generalisations. Often these conjectures will be false, and it is vital that we detect these cases and prune our search space appropriately.

There is a further, pragmatic argument for the inclusion of counterexample generation within the Grand Challenge. Great scientific problems, from landing a man on the moon to producing a computer to beat a grand master at chess, have been solved by making iterative improvements to a prototype and learning from failures. To apply this methodology to our Grand Challenge, we must encourage participation from industry, in order to ensure a supply of case study material, to get feedback on the tools we produce, and where appropriate, to try to influence software engineering practice. For this, we need to ensure that we are able to make a financial argument for the use of our tools even when they are at a prototype stage, and unable to deliver a fully verified end product. Being able to detect and present counterexamples that allow bugs to be identified is a way of ensuring some payback to our industrial partners. The current industrial preference for model checking over theorem proving can partly be explained by the ability of model checkers to present counterexample traces. Indeed, the recent interest in formal methods shown by the world's largest software company, Microsoft, can be traced back to the PRefix tool, presented to Microsoft by Intrinsa, a small Palo Alto start-up. This was a Lint-like tool for spotting possible bugs in source code. Microsoft were sufficiently impressed with the bugs detected to purchase the company, and their formal methods research has proceeded from there. It seems clear that the way to get our research into industrial use, and so to develop it to its full potential, is to develop tools which detect and handle counterexamples intelligently.

3 The Challenge of Counterexample Generation

The non-theorems that arise in program verification can sometimes be easy to find. They often occur close to the base case of a recursive data type, and evaluating the conjecture at some small values can be sufficient to detect the bug. However, there are a large number of cases where the counterexample is a far more subtle object, and it is here that the scientific challenge lies. An

example of a success story in this area is the application of formal methods to discovering attacks on security protocols. Here, the attacks are counterexamples to conjectures about the security of a system. The counterexamples may be quite large, for example up to 14 messages in [19], and may require an intruder to exploit quite subtle flaws in the protocol. Further challenges remain: for example, taking into account the mathematical properties of the cryptofunctions in use.

Other problems are well outside the scope of current techniques. For example, the book ‘Numerical recipes in C’ contains Knuth’s code for a (pseudo-)random number generator, [15, p. 283]. The code is designed to return a floating point number between 0 and 1. However, if certain very large seed values are used, the code can return a number outside of this range. With seed value 752 005 709¹, 13 out of the first 10,000 calls return a value very significantly outside the required range. The large seed value may seem ridiculous, but if you were seeding your generator on the number of seconds since 1 Jan 1970, this value would have occurred as a seed during 1993. The generation of these counterexamples remains well beyond the capabilities of current tools, and presents a substantial scientific challenge.

4 Survey of Counterexample Generation

Given how important the detection and presentation of counterexamples is to applications of formal methods, it is surprising how comparatively little attention it has received. A community of researchers interested in the subject is now beginning to emerge, with the third ‘Workshop on Disproving’ due to be held at IJCAR 2006².

Much past research has focused on fast generation of counterexamples to conjectures which are in some sense ‘obviously’ false. Model generators like MACE, [10], and FINDER, [18], can be used to enumerate finite domains to search for counterexamples. The Isabelle/HOL theorem prover, [12], now includes a tactic to search for small counterexamples, called quick-check, [5]. The Alloy modelling system generates counterexamples using a SAT-solver, [9]. For infinite data structures with recursive definitions, methods have been proposed by Protzen, [16], and Reif, [17]. Both can deal with small problems quite effectively, but are not suitable for large counterexamples, or for domains that cannot be easily enumerated. Model generation has also been proposed as a method for refuting non-theorems, [1,21].

Model checking can be a very effective method for finding counterexamples to false verification conditions, particularly in finite domains. As we have already remarked, the fact that model checking can produce counterexamples as well as provide guarantees is one reason for its increasing popularity in industry. Many researchers are now working on extending model checkers to non-finite domains, using ‘lazy’ and ‘on-the-fly’ techniques to construct the infinite models as they

¹ Knuth stated in the specification that the seed value can be any (large) number under 1 000 000 000. This is a known bug - the Numerical Recipes in C website contains a patch to fix the code.

² <http://www.cs.chalmers.se/~ahrendt/FLoC06-ws-disproving/>

are checked (e.g. [4]), with good results. Another large branch of current model checking research is in counterexample guided abstraction refinement, [6,3]. Here, the processing of counterexamples is used to guide management of the level of detail that is taken into account when attempting verification of a program, balancing tractability of the model checking problem against over-abstraction. Theorem provers can be used to check the feasibility of counterexample traces, [2]. A further current direction involves passing unproven conjectures from a theorem prover to a model checker to search for counterexamples, [14].

Our work in Edinburgh has lead to the development of the CORAL system³, which refutes incorrect inductive conjectures using a first-order theorem prover adapted to follow the ‘proof by consistency strategy’ of Comon and Nieuwenhuis, [7]. Its major successes have been in the field of protocol analysis, where it has been used to discover 6 previously unknown attacks on 3 group protocols, [19,20]. The protocols are modelled inductively following Paulson, [13], and the attacks found as counterexamples to security conjectures. CORAL proved to be particularly suitable for group protocols because they can be formalised very naturally in an inductive model. This is something that rival approaches, such as model checking, struggle with. In future, we plan to experiment with CORAL in other areas of formal verification.

Many of the big problems remain unsolved. For example, how to deal adequately with arithmetic, or how to explore very large or non-trivially enumerable spaces for counterexamples. Some current directions include trying to use more information from failed proof attempts to guide the counterexample search, [8], and to suggest patches for incorrect conjectures, [11].

5 Summary

We have argued that disproof and counterexample generation are vital areas for research in the development of practical program verification systems. There remain many exciting open problems, and milestones to pass, such as the automatic generation of counterexamples for Knuth’s random number bug, described above. Non-theorem detection therefore deserves to be included as a core theme in the Grand Challenge of a Program Verifier.

References

1. Ahrendt, W.: Deductive search for errors in free data type specifications using model generation. In: Voronkov, A. (ed.) CADE 2002. LNCS (LNAI), vol. 2392, Springer, Heidelberg (2002)
2. Ball, T., Cook, B., Lahiri, S., Zhang, L.: Zapato: Automatic theorem proving for predicate abstraction refinement. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 457–461. Springer, Heidelberg (2004)
3. Ball, T., Rajamani, S.: The SLAM toolkit. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, Springer, Heidelberg (2001)

³ <http://homepages.inf.ed.ac.uk/gsteel/coral>

4. Basin, D., Mödersheim, S., Viganò, L.: An on-the-fly model-checker for security protocol analysis. In: Proceedings of the, European Symposium on Research in Computer Security, pp. 253–270, 2003. Extended version available as Technical Report 404, ETH Zurich (2003)
5. Berghofer, S., Nipkow, T.: Random testing in Isabelle/HOL. In: 2nd International Conference on Software Engineering and Formal Methods (SEFM 2004), pp. 230–239 (2004)
6. Clarke, E., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. *Journal of the Association for Computing Machinery* 50(5), 752–794 (2003)
7. Comon, H., Nieuwenhuis, R.: Induction = I-Axiomatization + First-Order Consistency. *Information and Computation* 159(1-2), 151–186 (2000)
8. Dennis, L.A.: The use of proof planning critics to diagnose errors in the base cases of recursive programs. In: Basin, D., Rusinowitch, M. (eds.) *IJCAR 2004*. LNCS (LNAI), vol. 3097, pp. 47–58. Springer, Heidelberg (2004)
9. Jackson, D.: Alloy: a lightweight object modelling notation. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 11(2), 256–290 (2002)
10. McCune, W.: A Davis Putnam program and its application to finite first order model search. Technical report, Argonne National Laboratory (1994)
11. Monroy, R.: Predicate synthesis for correcting faulty conjectures: The proof planning paradigm. In: *Automated Software Engineering*, pp. 247–269 (2003)
12. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL. LNCS, vol. 2283. Springer, Heidelberg (2002)
13. Paulson, L.: The Inductive Approach to Verifying Cryptographic Protocols. *Journal of Computer Security* 6, 85–128 (1998)
14. Pike, L., Miner, P., Torres, W.: Model checking failed conjectures in theorem proving: a case study. Technical Report NASA/TM–2004–213278, NASA Langley Research Center (November 2004), <http://www.cs.indiana.edu/~lepik/pub-pages/unproven.html>
15. Press, W., Teukolsky, S., Vetterling, W., Flannery, B.: *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, Cambridge (1992)
16. Protzen, M.: Disproving conjectures. In: Kapur, D. (ed.) *11th Conference on Automated Deduction*, Saratoga Springs, NY, USA, June 1992. Springer Lecture Notes in Artificial Intelligence, vol. (607), pp. 340–354. Springer, Heidelberg (1992)
17. Reif, W., Schellhorn, G., Thums, A.: Flaw detection in formal specifications. In: Goré, R.P., Leitsch, A., Nipkow, T. (eds.) *IJCAR 2001*. LNCS (LNAI), vol. 2083, Springer, Heidelberg (2001)
18. Slaney, J.: FINDER: Finite Domain Enumerator. In: *Australian National University* (1995), <ftp://arp.anu.edu.au/pub/papers/slaney/finder/finder.ps.gz>
19. Steel, G., Bundy, A.: Attacking group multicast key management protocols using CORAL. *Electronic Notes in Theoretical Computer Science (ENTCS)* 125(1), 125–144 (2004) (Also available as Informatics Research Report EDI-INF-RR-0241. Presented at the ARSPA workshop 2004)
20. Steel, G., Bundy, A., Maidl, M.: Attacking a protocol for group key agreement by refuting incorrect inductive conjectures. In: Basin, D., Rusinowitch, M. (eds.) *IJCAR 2004*. LNCS (LNAI), vol. 3097, pp. 137–151. Springer, Heidelberg (2004)
21. Weber, T.: Bounded model generation for Isabelle/HOL. In: Basin, D., Rusinowitch, M. (eds.) *IJCAR 2004*. LNCS (LNAI), vol. 3097, pp. 27–36. Springer, Heidelberg (2004)