

## Goanna — A Static Model Checker

Ansgar Fehnker<sup>1</sup>, Ralf Huuck<sup>1</sup>, Patrick Jayet<sup>2\*</sup>, Michel Lussenburg<sup>2\*</sup>, and  
Felix Rauch<sup>1</sup>

<sup>1</sup> National ICT Australia Ltd. (NICTA)\*\* and University of New South Wales,  
Locked Bag 6016, NSW 1466, Australia

<sup>2</sup> Department of Computer Science, Swiss Federal Institute of Technology (ETH),  
CH-8092 Zurich, Switzerland

**Abstract.** In this work we present Goanna, the first tool that uses an off-the-shelf model checker for the static analysis of C/C++ source code. We outline its architecture and show how syntactic properties can be expressed in CTL. Once the properties have been defined the tool analyses source code automatically and efficiently. We demonstrate its applicability by presenting experimental results on analysing OpenSSL and the GNU coreutils.

### 1 Introduction

Formal design and analysis techniques are successfully applied to hardware. In fact, model checking parts of the chip design is common practice. However, the application of verification technology to existing and complex software has been much less successful.

The reasons are manifold: Full formal verification as done by interactive theorem proving is expensive. It requires a lot of time and expertise, making it often impractical for software that has a short life cycle, is not highly safety-critical, or suffers from a high pressure to market. Algorithmic verification techniques have to deal with software's infinite state space, requiring abstraction techniques to make properties of interest decidable. Suitable abstractions are typically hard to compute and the overall interaction required by the user in order to apply them to real-life software are often considerable.

One area that has been successful is static analysis [1, 2]. Approaches such as abstract interpretation, data flow analysis and other static checking techniques have made it into several industrial strength tools.

In this work we present Goanna [3], a static analysis tool for C/C++ source code based on model checking. It uses the NuSMV [4] model checker as the underlying verification engine, allows the specification of user defined properties

---

\* This work was carried out while visiting NICTA.

\*\* National ICT Australia is funded by the Australian Governments Department of Communications, Information Technology and the Arts and the Australian Research Council through Backing Australias Ability and the ICT Research Centre of Excellence programs.



use the weak until operator  $W$  here to include the second possibility. The latter can also point to unused variables, which is checked separately.

Our tool chain is depicted in Figure 1. We use `gcc` as a front end, as one of its features allows us to easily output the AST of C/C++ programs in an intermediate language. We parse the AST and on the one hand generate the CFG from it and on the other hand match patterns on the AST, which constitute the atomic propositions of a CTL formula expressing the desired property. We label the CFG with atomic propositions where their respective patterns were matched. Once the patterns and the CTL formula have been specified, the translation of the C/C++ source code into a suitable NuSMV model and its checking is fully automatic.

The current implementation is developed in OCaml. Goanna is easily integrated in Makefiles and, thus, is automatically supported by development environments such as Eclipse. A screen shot of Goanna running in combination with Eclipse can be found in Figure 2. As a result we obtain a seamless integration into the overall software development process.

### 3 Application

To evaluate the applicability of our tool, we examine two real-world open-source software packages: The GNU `coreutils`<sup>3</sup>, which provide basic file, shell and text manipulation utilities (59 kLoC<sup>4</sup>), and the `OpenSSL`<sup>5</sup> toolkit implementing the Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols (260 kLoC). We analyse the source code of these two packages on a recent 3.4 GHz Xeon-processor-based server.

Analysing the whole source with our current Goanna tool (which has not yet been optimised) takes slightly less than 2 minutes for the `coreutils` and slightly less than 29 minutes for `OpenSSL`. The latter is somewhat distorted by a single pathological file that takes almost 12 minutes to analyse. In practice, analysis times are typically much shorter, because the analysis can be done incrementally on the set of recently changed files only and a more in-depth study of Goanna's analysis times shows that a large majority of source files is analysed quite quickly. In fact, 72% of all source files in the `coreutils` are analysed in less than 1 second and 95% under 5 seconds. Similarly, for `OpenSSL` 83% of all files are analysed in under 1 second and again 95% under 5 seconds.

Note that the current prototype has not yet been optimised regarding execution time. Hence, there is still a lot of room for performance improvements, for example by optimising the search in the AST for patterns of interest (which currently contains redundant searches for different properties), the OCaml library we use to conduct the search on the AST (which is convenient to use but not efficient), or by changing the way in which we use NuSMV (which is cur-

---

<sup>3</sup> <http://www.gnu.org/software/coreutils/>

<sup>4</sup> LoC = Lines of Code, kLoC = 1000 Lines of Codes

<sup>5</sup> <http://www.openssl.org/>

rently invoked with rather large chunks of source code at the time that could be reduced to smaller pieces).

Looking at the memory requirements of our tool we find that the maximum memory consumption of the analysis is about 65 MiB to analyse the coreutils and about 113 MiB for OpenSSL respectively. This is in both cases much below the limit set by today's PCs used by developers.

The above numbers show that the tool is already quite usable in practice. A full evaluation of course requires also an analysis of the precision of the tool, with looking at the number of real bugs found and the number of false positives reported. Such a study is very time consuming and we are still in the process of qualitatively evaluating Goanna regarding its precision. Preliminary results indicate that the precision of our approach is comparable to standard static analysis.

## 4 Conclusion

In this work we presented Goanna, the first static analyser purely based on an off-the-shelf model checker. We demonstrated that the approach scales well to real-life software, making it suitable for the integration into the overall software development process.

While Goanna is fast, it is not yet more precise than traditional static analysis. However, we anticipate to improve on this by incorporating more semantic-based software model checking techniques such as predicate abstraction [6]. The foundation of this integration has been laid by having a uniform framework for static analysis as well as traditional model checking.

## References

1. Engler, D.R., Musuvathi, M.: Static analysis versus software model checking for bug finding. In: "VMCAI '04: 5th Intl. Conference Verification, Model Checking and Abstract Interpretation". (2004) 191–210
2. Blanchet, B., Cousot, P., Cousot, R., Feret, J., Mauborgne, L., Miné, A., Monniaux, D., Rival, X.: Design and implementation of a special-purpose static program analyzer for safety-critical real-time embedded software. (2002) 85–108
3. NICTA: The Goanna Project. (<http://ertos.nicta.com.au/research/goanna/>)
4. Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A.: NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking. In: Proc. International Conference on Computer-Aided Verification (CAV 2002). Volume 2404 of LNCS., Springer (2002)
5. Schmidt, D.A., Steffen, B.: Program analysis as model checking of abstract interpretations. In: SAS '98: Proceedings of the 5th International Symposium on Static Analysis, London, UK, Springer-Verlag (1998) 351–380
6. Ball, T., Cook, B., Levin, V., Rajamani, S.K.: SLAM and Static Driver Verifier: Technology transfer of formal methods inside Microsoft. In: IFM 2004: 4th International Conference on Integrated Formal Methods. Volume 2999 of LNCS., Springer-Verlag (2004) 1–20