

Analysing Embedded System Software

–Extended Abstract–

Ansgar Fehnker, Ralf Huuck, Felix Rauch, and Sean Seefried

National ICT Australia Ltd. (NICTA)* and University of New South Wales,
Locked Bag 6016, Sydney NSW 1466, Australia

Abstract. The verification of real-life C/C++ code is inherently hard. Not only are there numerous challenging language constructs, but the precise semantics is often elusive or at best vague. This is even more true for systems software where non-ANSI compliant constructs are used, hardware is addressed directly and assembly code is embedded. In this work we present a lightweight solution to detect software bugs in C/C++ code. Our approach performs static checking on C/C++ code by means of model checking. While it cannot guarantee full functional correctness, it can be a valuable tool to increase the reliability and trustworthiness of real-life system code. This paper explains the general concepts of our approach, discusses its implementation in our C/C++ checking tool *Goanna*, and presents some performance results on large software packages.

1 Introduction

Showing the full functional correctness of system software written, e.g., in C/C++, is a major challenge. It requires a precise understanding of the underlying semantics, typically needs to include an abstract hardware model, and has to give a full functional proof. There are a number of projects currently undertaking this task [1–3] supported by interactive theorem provers. While this is the only way to guarantee the full correctness of a program, it requires substantial resources both in time as well as in the number of highly qualified people.

On the other hand, commercial system software has a high pressure to market, needs to run on various platforms and is rewritten frequently, making the above approach even more challenging. There are a number of lightweight analysis approaches that seek to complement full verification by detecting software bugs at the coding stage and, thus, increasing the reliability and trustworthiness of the code. Those tools make a limited but practical contribution to program correctness and can support full verification by reducing property violations in early stages.

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

The model-checking community has made significant advances in recent years to cover realistic C/C++ programs and produced a number of powerful tools [4–7]. However, they are not yet well-suited for real-life embedded system code [8, 9]. On the other hand, commercial static analysis tools [10–13] cope well with most C/C++ code and make a valuable contribution to software correctness. In contrast to model checking tools, static analysers typically do not allow for any user-defined specifications, but rather implement a set of independent analysis heuristics or allow specification which are less expressive than the temporal logics used by model checkers.

In this work we present a static analysis approach based on model checking. While we retain the flexibility and power of temporal logics specifications, we are able to handle any parsed C/C++ code in a uniform manner. In particular, we present the underlying idea of translating C/C++ checks into model checking properties, which can then be checked by one single analyzer, instead of a set of static analysis heuristics. In our case we use the NuSMV [14] model checker as back end. Moreover, we present some implementation details of our checker *Goanna* and its performance on the source of large, real-life open-source software packages.

Section 2 describes our underlying framework, while Section 3 presents some of our performance results and Section 4 discusses the current state of our research as well as ongoing and future work.

2 Static Analysis by Means of Model Checking

In this section we describe how to statically check properties of C/C++ source code by means of model checking. This approach has been inspired by [15, 16] and is also followed by [17, 18].

Using a model checker for solving static analysis problems has a number of advantages. All properties can be expressed in a single, flexible analysis engine. This means that it is easy to add new checks by adding new checking properties. In addition, the analysis scales well with increasing number of properties. The details of our path-sensitive, intra-procedural analysis can be found in [19].

The basic idea is to annotate the control flow graph (CFG) of a program with atomic propositions of interest. In order to check, e.g., for uninitialised variables, we can identify atomic propositions $decl_q$, $read_q$ and $write_q$, representing program locations where a variable q is declared but not initialised, where it is read from or written to, respectively, and mark those locations in the CFG accordingly. The atomic propositions are identified by purely syntactic criteria on the abstract syntax tree (AST) of the program by means of a pattern language. We define patterns for each proposition, e.g., a variable is written to if it occurs on the left hand side of an assignment statement and so on. Once identified, the proposition is placed on the node in the CFG most closely corresponding to the nodes in the AST where it was identified.

An example of the resulting annotated CFG can be found in Figure 1. This representation is already very close to a Kripke structure and we can model check




Fig. 1. Annotated CFG

that structure for properties of interest. For instance, checking for uninitialised variables can be expressed in CTL as:

$$AG \text{ decl}_q \Rightarrow (A \neg \text{read}_q \ W \ \text{write}_q)$$

This means we require that on all program paths if a variable q is declared it must not be read until it has been written or it will not be written at all. We use the *weak until* operator W here to include the second possibility. The latter can also point to unused variables, which is checked separately.

In the same style we can express other properties on correct pointer handling, variable usage or memory allocation and deallocation. Moreover, it allows specifying application specific properties to handle general programming guidelines, API-specific rules or even hardware/software interface rules for device drivers.

Once the patterns relevant for matching atomic propositions have been defined and the CFG has been annotated, it is straightforward to translate the annotated graph automatically into the input format of a model checker. Adding new checks only requires one to define the property to be checked and the patterns representing atomic propositions. All other steps can be fully automated.

Although this framework was developed in first instance for C/C++ it can be also extended to deal with embedded assembly code. This is important for the embedded systems space, since interaction with the hardware is frequently implemented as embedded assembly code. In particular, we take C/C++ and ARMv6 assembly interface information for our analysis into account, check for compliance of embedded assembly code with its C/C++ interface, and perform various checks on the pure assembly level. The combined analysis of C/C++ code with embedded assembly code enhances, in addition, the precision of the analysis.

3 Implementation and Evaluation

The aforementioned approach has been implemented in our program analyser Goanna, using the open source model checker NuSMV [14] as a generic back-end analysis engine. The surrounding code for pattern matching structures of interest, property definitions, CFG generation, translation into NuSMV, and representation of analysis results is written in OCaml. Moreover, Goanna can be invoked just like the gcc/g++ compiler and, therefore, integrates seamlessly into standard development environments such as Eclipse (cf. Figure 2).

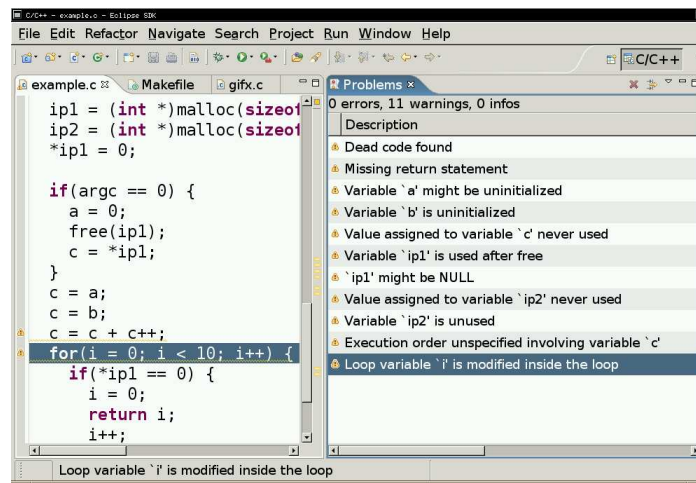


Fig. 2. Eclipse integration of Goanna

We evaluated Goanna on a number of open source packages ranging from highly optimized system software such as the L4 microkernel¹ to large application code bases such as the 260 kLoC² OpenSSL package.

For an unoptimized version of Goanna some run-time results for OpenSSL are shown in Figure 3. It shows that over 80% of all files are analyzed within 1 second and that 99% of all files are analyzed within 5 seconds. The whole analysis takes less than 15 minutes. Proportionally, the time spent purely in NuSMV is mostly negligible with 98.7% of all files being analyzed in less than 2 seconds.

The run-times of Figure 3 are based on checking for 15 properties ranging from simple uninitialized variables, over potential null-pointer dereferences, to memory leaks. It is worth to mention that increasing the number of properties typically scales well in our framework as it only increases the number of labels and property specifications in the same NuSMV model, which is handled well by

¹ <http://14hq.org/>

² LoC = Lines of Code

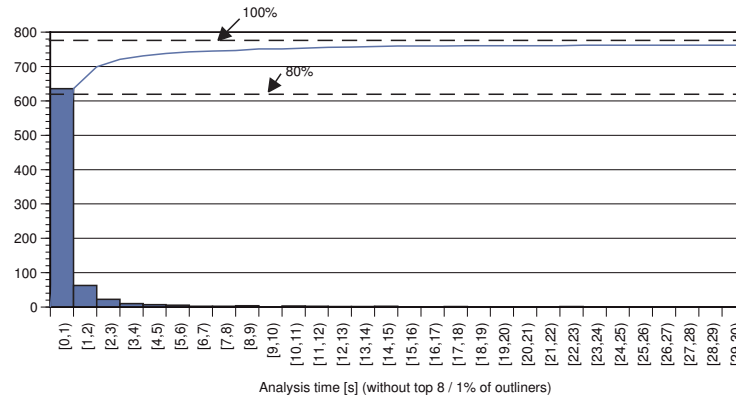


Fig. 3. Run-times for Goanna analysis on OpenSSL source files

the model checker. For instance, increasing the number of properties from one to 15 only doubled the overall analysis time.

Moreover, we found that the analysis time is not only well within the same order of magnitude as the compile time, but that the memory requirements of the analysis fit easily in the RAM of current developer machines.

The analysis of C/C++ code with embedded assembly code was evaluated for Pistachio 0.4 implementation³ of L4, compiling for an ARM SA1100 architecture. It contains 54 C++ files, two of which have embedded assembly blocks (3.7%), and they include a total of 72 header files, of which 10 have embedded assembly blocks (13.8%). The additional assembly analysis lead to a modest increase from 75.9 seconds to 77.3 seconds, which is an increase of only 1.4 seconds or 1.8%.

4 Conclusion

Summary. In this work we presented our framework and results on model checking system software by means of static analysis. We showed how to easily encode static checks as model checking properties, providing the basis for an extendable and flexible checker. Moreover, we implemented our analysis framework in Goanna, the first static checker using NuSMV as its analysis engine, and presented some run-time and scalability results. We showed that this is a viable solution that can be integrated well in the software development process.

Ongoing and Future Work. Currently, we are working on improving the precision of the analysis. Future work will focus on further increasing the performance of Goanna, integrating a full inter-procedural analysis and defining a user interface for property specification.

³ <http://14hq.org/>

Acknowledgements We thank Bernard Blackham, Jörg Brauer, Patrick Jayet and Michel Lussenburg for their implementation efforts and general contributions.

References

1. Tuch, H., Klein, G., Heiser, G.: OS verification — now!, Santa Fe, NM, USA (2005) 7–12
2. Hohmuth, M., Tews, H.: The VFiasco approach for a verified operating system. In: Proc. 2nd ECOOP Workshop on Programm Languages and Operating Systems, Glasgow, UK (2005)
3. Gargano, M., Hillebrand, M., Leinenbach, D., Paul, W.: On the correctness of operating system kernels. In: Proc. 18th International Conference on Theorem Proving in Higher Order Logics (TPHOLs’05), Oxford, UK (2005) 1–16
4. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL. (2004) 232–244
5. Ball, T., Rajamani, S.K.: The SLAM Toolkit. In: Intl. Conf. on Computer Aided Verification (CAV ’01), London, UK, Springer-Verlag (2001) 260–264
6. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2004). Volume 2988 of LNCS., Springer (2004) 168–176
7. Chaki, S., Clarke, E.M., Groce, A., Strichman, O.: Predicate Abstraction with Minimum Predicates. In: CHARME. (2003) 19–34
8. Schlich, B., Kowalewski, S.: Model checking C source code for embedded systems. In: Proc. of the IEEE/NASA Workshop on Leveraging Applications of Formal Methods, Verification, and Validation, NASA/CP-2005-212788 (2005)
9. Mühlberg, J., Lüttgen, G.: BLASTing Linux code. In: Proc. of the 11th International Workshop on Formal Methods for Industrial Critical Systems (FMICS 06). Volume 4346 of LNCS., Springer-Verlag (To appear)
10. Coverity: Prevent for C and C++. (<http://www.coverity.com>)
11. Gimpel Software: Flexelint for C/C++. (<http://www.gimpel.com/html/flex.htm>)
12. Klocwork: K7. (<http://www.klocwork.com/products/klocworkk7.asp>)
13. Microsoft: Prefast. (<http://www.microsoft.com/whdc/devtools/tools/PREfast.mspcx>)
14. 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: CAV (2002). LNCS 2404, Springer (2002)
15. Schmidt, D.A., Steffen, B.: Program analysis as model checking of abstract interpretations. In: Intl. Symposium on Static Analysis (SAS ’98), London, UK, Springer-Verlag (1998) 351–380
16. Schmidt, D.A.: Data flow analysis is model checking of abstract interpretations. In: ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL ’98), New York, NY, USA, ACM Press (1998)
17. Holzmann, G.: Static source code checking for user-defined properties, Pasadena, CA, USA (2002)
18. Dams, D., Namjoshi, K.S.: Orion: High-precision methods for static error analysis of C and C++ programs. Bell Labs Technical Memorandum ITD-04-45263Z, Lucent Technologies (2004)
19. Fehnker, A., Huuck, R., Jayet, P., Lussenburg, M., Rauch, F.: Model checking software at compile time. In: Proc. of the 1st IEEE & IFIP International Symposium on Theoretical Aspects of Software Engineering (TASE), Shanghai, China (2007, to appear)