

Skink: Static Analysis of Programs in LLVM Intermediate Representation (Competition Contribution)

Franck Cassez^(✉), Anthony M. Sloane, Matthew Roberts, Matthew Pigram,
Pongsak Suvanpong, and Pablo Gonzalez de Aledo

Macquarie University, Sydney, Australia
franck.cassez@mq.edu.au

Abstract. SKINK is a static analysis tool that analyses the LLVM intermediate representation (LLVM-IR) of a program source code. The analysis consists of checking whether there is a feasible execution that can reach a designated error block in the LLVM-IR. The result of a program analysis is “correct” if the error block is not reachable, “incorrect” if the error block is reachable, or “inconclusive” if the status of the program could not be determined. In this paper, we introduce SKINK 2.0 to analyse single and multi-threaded C programs.

1 Overview

SKINK is a static analysis tool that analyses the LLVM intermediate representation (LLVM-IR) of a source program. For instance, SKINK can analyse C/C++ programs using the LLVM-IR as generated by the CLANG compiler. The objective of the static analysis is to check whether a program is correct w.r.t. a given specification. For C/C++ programs, the specification is provided via `assert(condition)` statements in the C program. The aim of the analysis is to determine whether a condition can be violated. In the SV-COMP setting, `assert` calls `__VERIFIER_error` if the condition is false.

The LLVM-IR representation consists of a collection of *functions* made up of *blocks*. A block represents a sequence of simple instructions (e.g., `store`, `load`, function calls) and ends with a terminating instruction such as a branch that points to the next block(s). The LLVM-IR we analyse contains a designated “error” block that corresponds to a call to `__VERIFIER_error`. A (feasible) program trace (execution) that contains the error block is an *error trace*. A program is incorrect if and only if it can generate an error trace.

2 Verification Approach

SKINK’s strategy to determine whether an error trace exists uses the iterative *refinement of the trace abstraction* algorithm of Heizmann [1, 2]. First the LLVM-IR of a source program is mapped to a control flow graph (CFG) which is a

finite labeled automaton. The labels are the “basic blocks” and the “choices” (branching) of the LLVM-IR. In the automaton the labels are letters and do not carry any special meaning.

The (regular) language accepted by the CFG is the set of traces *leading to an error block*. These traces are *abstract error traces*: the CFG does not give any semantics to the labels, and it is not guaranteed that any such trace is actually *feasible* in the concrete program.

Checking whether a program is correct reduces to determining whether an abstract error trace is *feasible* or equivalently whether the language of the CFG contains a *feasible* abstract error trace. To determine whether a trace is feasible, we take into account the semantics of the instructions of the basic blocks. This is achieved by encoding a trace of the CFG as a logical statement and checking whether this statement is satisfiable or not. If satisfiable, a feasible error trace has been found and the program is incorrect. Otherwise, if a trace t is spurious, an *interpolant automaton* can be computed that accepts t and other traces that are infeasible for the same reason as t [1, 2]. In the latter case, we can *refine* the CFG and look for an error trace in the language of the CFG minus the language accepted by the interpolant automaton. When this iterative refinement process stops¹ either no error traces remain and the program is correct or a feasible error trace is discovered and the program is incorrect.

This algorithm can also be used for checking *concurrent* programs using the *product of the CFGs* in each thread. This is in general not very effective as the number of interleavings grows exponentially in the number of threads. The algorithm implemented in SKINK 2.0 to analyse multi-threaded C programs extends our previous work that showed how to combine trace abstraction refinement and *partial-order reduction* [3].² In SKINK 2.0 we have combined trace abstraction refinement with state-of-the-art *dynamic partial order reduction* techniques [4].

3 Software Architecture

SKINK 2.0 is developed in SCALA and can directly analyse LLVM-IR programs. SKINK 2.0 is currently able to analyse programs in C (via CLANG) but it is trivially expandable to any language that can be compiled to LLVM-IR.

Front-end: SKINK’s front-end is written using our SBT-RATS parser generator [5]³ and our KIAMA SCALA library for language processing [6].⁴ We have developed a Scala-LLVM parser that can read LLVM-IR and build an abstract syntax tree (AST). Semantic analysis is performed on the AST to recover information such as variable types. SKINK 2.0 constructs CFGs for functions from the AST using the KIAMA attribute grammar methods.

¹ It may never stop and in this case the analysis is inconclusive.

² The implementation introduced in [3] has nothing in common with SKINK; it was limited to analysing programs written in a custom input language (but not C) and implemented static partial-order reductions algorithms.

³ <https://bitbucket.org/inkytonik/sbt-rats>.

⁴ <https://bitbucket.org/inkytonik/kiama>.

Middle-end: Our SCALA library AUTOMAT⁵ provides the automata-theoretic operations (union, intersection, DFS, partial order reduction) that are needed in the refinement algorithm. This is used to obtain candidate abstract error traces (via a test for language emptiness) and to construct the refinements (difference between two regular languages). On top of the automata-based refinement algorithm, SKINK 2.0 provides two core functionalities. The first is the encoding of an abstract error trace into an SSA form and eventually a logical formula; this logical formula is satisfiable if and only if the trace is feasible. Satisfiability is determined by an SMT-solver (see *Back-end* section below). The second is the computation of an *interpolant automaton* which is based on an annotation of an infeasible trace with *invariants*.

Back-end: To check whether a (symbolic) abstract error trace is feasible we use a SCALA abstraction over the SMTLIB standard for common languages and interfaces for SMT solvers. Our library MQ-SCALA-SMTLIB⁶ provides this abstraction. MQ-SCALA-SMTLIB was also developed using SBT-RATS and KIAMA.

We support most of the SMT-solvers (including Z3⁷, SMTINTERPOL⁸ and CVC4⁹) via a common SCALA abstract interface. As a result we can choose which solver to use at run time, and we may use multiple and different solvers during the same program analysis. In the current implementation, SKINK 2.0, we mostly use Z3, SMTInterpol and CVC4, depending on the theories and operations we need (linear integer arithmetic, arrays, bitvectors, interpolants) and on the SV-COMP categories to analyse.

4 Strengths and Weaknesses

SKINK 2.0 does not support the full LLVM-IR assembly language and our front-end parser may fail to parse some LLVM-IR input. Another limitation of SKINK 2.0 is that we use the LLVM *inlining* capability (`opt -inline`) to obtain a single CFG for each LLVM-IR. This may fail preventing the subsequent program analysis. These limitations should be overcome in the next months by extending our front-end Scala-LLVM parser and implementing our modular analysis technique [7]. We may assume unbounded integers in the analysis and this may result in false negatives due to overflow/underflow errors (in our tests it happened once in the ControlFlow category).

On the positive side, SKINK 2.0 can analyse programs that can be compiled into LLVM-IR which makes it usable on a variety of languages including C/C++, Objective C and Swift. A major strength of SKINK 2.0 is that it can discover loop invariants (interpolants) and is able to establish program correctness.

⁵ <https://bitbucket.org/franck44/automat>.

⁶ <https://bitbucket.org/franck44/mq-scala-smtlib>.

⁷ <https://github.com/Z3Prover/z3>.

⁸ <https://ultimate.informatik.uni-freiburg.de/smtinterpol/>.

⁹ <http://cvc4.cs.nyu.edu/web/>.

Our abstract SCALA solver library MQ-SCALA-SMTLIB provides access to a number of theories (Arrays, BitVectors) and solver capabilities (generate interpolants). SKINK 2.0 is, to the best of our knowledge, the only tool that combines trace abstraction refinement with a version (source-DPOR) of the optimal state-of-the-art *dynamic partial order reduction* algorithm [4]. This enables us to efficiently verify some programs in the Concurrency benchmarks category.

5 Set up and Configuration

Participation Statement: SKINK opts-out from all categories except *Integer and Control Flow*, *Concurrency* and *BitVectors*.

Set up and Configuration: SKINK 2.0 is available from <http://science.mq.edu.au/~fcassez/sw/skinkv2.0.tgz>. The archive includes all dependencies needed to run it on Ubuntu Xenial Xerus 64-bit (16.04.1). `skink.sh` is the simplest and the recommended way to run this SKINK 2.0 distribution¹⁰. `skink.sh` should be passed the C file on which analysis is to be performed. It will place along that file the verification output (`.verif`) and the witness file (`.graphml`) as appropriate.

6 Software Project and Contributors

SKINK 2.0 is developed by F. Cassez, A. M. Sloane, M. Roberts, M. Pigram, P. Gonzalez, P. Suvanpong at the Department of Computing, Macquarie University. The libraries used in SKINK 2.0 are open-source software. More information can be found at <http://science.mq.edu.au/~fcassez/software-verif.html>.

References

1. Heizmann, M., Hoenicke, J., Podelski, A.: Refinement of trace abstraction. In: Palsberg, J., Su, Z. (eds.) SAS 2009. LNCS, vol. 5673, pp. 69–85. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-03237-0_7](https://doi.org/10.1007/978-3-642-03237-0_7)
2. Heizmann, M., Hoenicke, J., Podelski, A.: Software model checking for people who love automata. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 36–52. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-39799-8_2](https://doi.org/10.1007/978-3-642-39799-8_2)
3. Cassez, F., Ziegler, F.: Verification of concurrent programs using trace abstraction refinement. In: Davis, M., Fehnker, A., McIver, A., Voronkov, A. (eds.) LPAR 2015. LNCS, vol. 9450, pp. 233–248. Springer, Heidelberg (2015). doi:[10.1007/978-3-662-48899-7_17](https://doi.org/10.1007/978-3-662-48899-7_17)
4. Abdulla, P.A., Aronis, S., Jonsson, B., Sagonas, K.F.: Optimal dynamic partial order reduction. In: Jagannathan, S., Sewell, P. (eds.) POPL 2014, San Diego, CA, USA, 20–21 January 2014, pp. 373–384. ACM (2014)
5. Sloane, A.M., Cassez, F., Buckley, S.: The sbt-rats parser generator plugin for ala (tool paper). In: SCALA 2016, pp. 110–113. ACM, New York (2016)

¹⁰ The required Scala run-time and libraries are bundled into the `skink.jar` file.

6. Sloane, A.M.: Lightweight language processing in Kiama. In: Fernandes, J.M., Lämmel, R., Visser, J., Saraiva, J. (eds.) GTTSE 2009. LNCS, vol. 6491, pp. 408–425. Springer, Heidelberg (2011). doi:[10.1007/978-3-642-18023-1_12](https://doi.org/10.1007/978-3-642-18023-1_12)
7. Cassez, F., Müller, C., Burnett, K.: Summary-based inter-procedural analysis via modular trace refinement. In: FSTTCS 2014, LIPIcs, vol. 29, pp. 545–556. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2014)