

# WOLVERINE: Battling Bugs with Interpolants<sup>\*</sup>

## (Competition Contribution)

Georg Weissenbacher<sup>1</sup>, Daniel Kroening<sup>2</sup>, and Sharad Malik<sup>1</sup>

<sup>1</sup> Department of Electrical Engineering, Princeton University

<sup>2</sup> Department of Computer Science, Oxford University

**Abstract.** WOLVERINE is a software verifier that checks safety properties of sequential ANSI-C and C++ programs, deploying Craig interpolation to derive program invariants. We describe the underlying approach and the architecture, and provide instructions for installation and usage.

## 1 Approach

WOLVERINE [8] is a software verification tool for ANSI-C and C++ programs that aims at finding either a Hoare-style correctness proof or a counterexample for a given reachability property. The tool is an implementation of the interpolation-based lazy abstraction algorithm [9] outlined in Figure 1. A description of the steps ① to ⑤ of Figure 1 is provided below.

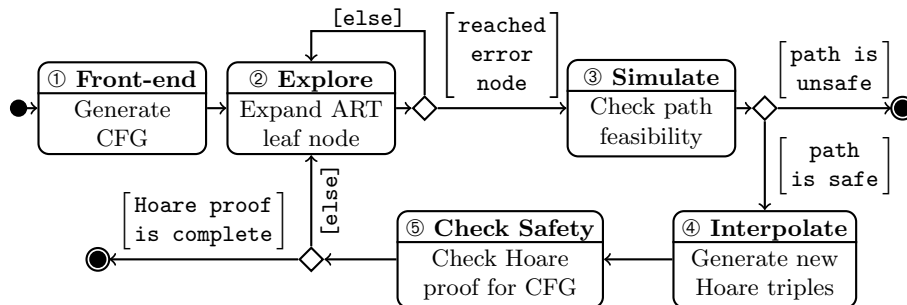


Fig. 1: UML activity diagram describing the work-flow of WOLVERINE

- ① WOLVERINE generates a control flow graph (CFG) representation of the program and encodes reachability properties using assertions/error nodes.
- ② Following the lazy abstraction paradigm established by [5], WOLVERINE constructs an abstract reachability tree (ART). To this end, it explores the paths of the CFG (in a depth first search manner) until it encounters an assertion.<sup>3</sup>

Corresponding author: Georg Weissenbacher (georg.weissenbacher@magd.oxon.org)

<sup>\*</sup> Supported by a gift from the Intel Labs Academic Research Office.

<sup>3</sup> The search algorithm of WOLVERINE 0.5c incorporates a constant propagation domain in order to enable early pruning of infeasible execution traces.

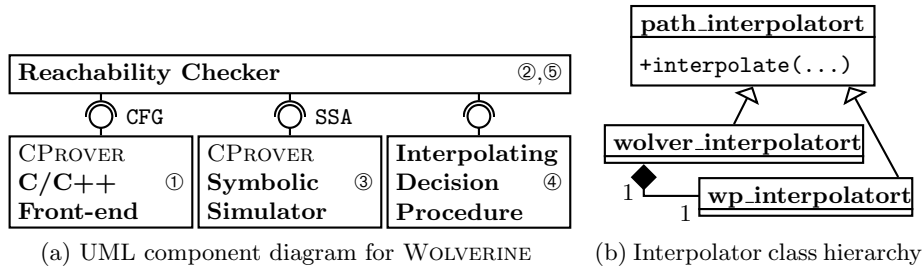


Fig. 2: Software architecture of WOLVERINE

- ③ Given a path that reaches an assertion, WOLVERINE deploys symbolic simulation to determine whether the path corresponds to a feasible program execution violating the assertion; such unsafe executions are reported.
- ④ If the path is safe, WOLVERINE uses Craig interpolation to generate Hoare triples that prove that the assertion holds (c.f. [4]) and updates the edges and nodes of the ART accordingly: the spurious counterexample serves as a catalyst for refining the current approximation of safely reachable states [2].
- ⑤ If the Hoare triples of the ART are sufficient to prove the safety of all paths of the CFG, WOLVERINE concludes that the program is correct. Otherwise, the tool continues to expand paths that are not yet covered (step ②).

## 2 Software Architecture

Figure 2a shows the components and architecture of WOLVERINE. Our implementation uses the front-end (①) and the symbolic simulator (③) of the CPROVER framework (<http://www.cprover.org>). WOLVERINE uses an interpolating decision procedure (④) to extract Hoare triples from infeasible paths. To this end, the tool deploys its built-in decision procedure for equality logic with uninterpreted functions and limited support for bit-vectors [3, 6, 7] and falls back on the weakest precondition should this interpolator fail (see Figure 2b).

## 3 Tool Setup and Usage

*Installation.* Binaries for Linux, Windows, and MacOS X can be downloaded from the project website (<http://www.cprover.org/wolverine>) and should be deployed in a directory listed in the `PATH` environment variable. WOLVERINE requires a pre-processor (`c1.exe`, which is part of Visual Studio Express, on Windows and GNU’s `gcc` on Unix-based platforms) and the header files typically packaged with it to be installed.

*Usage.* WOLVERINE must be executed from within the Visual Studio command prompt on Windows or a terminal on Linux and Mac OS X, and accepts options and source file names of the program to be verified as operands. By default, WOLVERINE scans the program for assertions and checks whether they hold. If

executed with the option `--error-label ERROR`, WOLVERINE checks whether the label `ERROR` is reachable.

By default, WOLVERINE assumes that the host platform and the target platform for the program under test coincide. Therefore, in order to verify a Windows device driver on a Linux host, the options `--no-library --i386-win32` are recommended (but were not applied in the competition). The target processor architecture of the program under test has to be specified using the options `--32` or `--64` where it differs from the host. In the competition, these options were applied accordingly to all benchmarks.

## 4 Strengths and Limitations

While WOLVERINE shares many of the characteristics of predicate abstraction-based verifiers (most prominently, SLAM [1]), it avoids the computationally expensive image computation required to construct the abstraction (c.f. [9]), enabling the rapid detection of counterexamples (discussed in [8]).

WOLVERINE’s performance is contingent on the Hoare triples that the interpolating decision procedure derives from spurious counterexamples. The inherent properties of interpolants typically enable concise abstractions. A “diverging” sequence of predicates, however, can result in a failed verification attempt. The built-in interpolator of WOLVERINE version 0.5c provides no support for linear arithmetic, quantified invariants, and heap models. In the competition, this led to a sub-optimal performance of WOLVERINE for benchmarks containing arithmetic expressions, unbounded arrays, or dynamic data structures. Moreover, WOLVERINE does currently not support the verification of concurrent programs.

## References

1. Ball, T., Rajamani, S.K.: The SLAM project: Debugging system software via static analysis. In: POPL, pp. 1–3. ACM (2002)
2. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: CAV. LNCS, vol. 1855, pp. 154–169. Springer (2000)
3. D’Silva, V., Purandare, M., Weissenbacher, G., Kroening, D.: Interpolant strength. In: VMCAI. LNCS, vol. 5944, pp. 129–145. Springer (2010)
4. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL, pp. 232–244. ACM (2004)
5. Henzinger, T.A., Jhala, R., Majumdar, R., Sutre, G.: Lazy abstraction. In: POPL, pp. 58–70. ACM (2002)
6. Kroening, D., Weissenbacher, G.: Lifting propositional interpolants to the word-level. In: FMCAD. pp. 85–89. IEEE (2007)
7. Kroening, D., Weissenbacher, G.: An interpolating decision procedure for transitive relations with uninterpreted functions. In: HVC. LNCS, vol. 6405. Springer (2011)
8. Kroening, D., Weissenbacher, G.: Interpolation-based software verification with Wolverine. In: CAV, LNCS, vol. 6806, pp. 573–578. Springer (2011)
9. McMillan, K.L.: Lazy abstraction with interpolants. In: CAV, LNCS, vol. 4144, pp. 123–136. Springer (2006)