



# The Static Analyzer Infer in SV-COMP (Competition Contribution)

Matthias Kettl( ) and Thomas Lemberger()

LMU Munich, Germany

**Abstract.** We present `INFER-SV`, a wrapper that adapts `INFER` for SV-COMP. `INFER` is a static-analysis tool for C and other languages, developed by Facebook and used by multiple large companies. It is strongly aimed at industry and the internal use at Facebook. Despite its popularity, there are no reported numbers on its precision and efficiency. With `INFER-SV`, we take a first step towards an objective comparison of `INFER` with other SV-COMP participants from academia and industry.

## 1 Facebook Infer

`INFER` [6] is a compositional and incremental static-analysis tool developed at Facebook. `INFER` supports a wide array of analyses; this includes memory safety, buffer overruns, performance constraints and different reachability analyses for C, C++, Objective C, Java, C#, and .Net. For memory analysis, `INFER` uses bi-abduction [7] with separation logic [14]. `INFER` supports the integration of new abstract domains through the abstract-interpretation framework `Infer:AI`. `INFER` analyzes programs compositionally (building method summaries) and incrementally (only analyzing changed program parts). In contrast to most other tools that participate in SV-COMP, `INFER` is not an academic verifier. Instead, it is aimed at practical use during software development. This has direct implications on the development focus: When `INFER` is told to incrementally analyze software, it outputs only newly discovered bugs and does not re-report bugs found in previous analyses. This allows developers to ignore warnings not deemed relevant and reduces the cognitive burden on developers due to false alarms. Multiple large companies use `INFER`—among others: Amazon Web Services, Facebook, Microsoft, Mozilla, and Spotify. At the time of this writing, `INFER` has more than 12 000 stars on GitHub and was forked over 1 500 times. Despite its popularity, there are no reported numbers on `INFER`'s precision and soundness. With the participation of `INFER` in the C language track of SV-COMP '22, we hope to take a first step towards an objective comparison of `INFER` with other verifiers.

The following other commercial verifiers participate in SV-COMP '22: `2LS` [16], `CBMC` [10], `CRUX`<sup>1</sup>, `FRAMA-C` [5], `VERIABS` [12], and `VERIFUZZ` [9].

<sup>1</sup> <https://crux.galois.com/>

## 2 Infer in SV-COMP

### 2.1 Infer-SV

**Verification.** We provide the wrapper `INFER-SV` to adapt `INFER` to the SV-COMP specification format for program properties. `INFER-SV` parses the property to analyze, adjusts the program under analysis for `INFER`, runs `INFER` with fitting analyses, and reports a verification verdict based on the feedback produced by `INFER`. `INFER-SV` supports the following SV-COMP program properties:

*no-overflow.* The aim is to check for arithmetic overflows on signed-integer types. `INFER-SV` runs `INFER`'s buffer-overflow analysis<sup>2</sup> to detect these.

*unreach-call.* The aim is to check for reachable calls to function `reach_error`. `INFER` provides a function-call reachability analysis<sup>3</sup>, but this analysis proved very imprecise. To mitigate this, `INFER-SV` performs a program transformation<sup>4</sup>: It replaces each call to function `reach_error` with an overflow-provoking statement `int __reach_error_x = 0x7fffffff + 1`. No task with property `unreach-call` contains a signed-integer overflow, so the original reachability property holds if and only if any of the introduced overflows is reachable. `INFER-SV` runs `INFER`'s buffer-overflow analysis on the transformed program to check this.

*valid-memsafty.* The aim is to check for invalid pointer dereferences, invalid frees of memory, and memory leaks. To analyze memory safety, `INFER-SV` uses two analyses: bi-abduction<sup>5</sup> and `Infer:Pulse`<sup>6</sup>. SV-COMP requires verifiers to report the concrete type of violation detected: *valid-deref*, *valid-memtrack*, or *valid-free*. `INFER-SV` analyzes the error codes reported by `INFER` to determine the exact violation found. If `INFER` reports multiple fitting warnings, we take the first.

**Witnesses.** SV-COMP requires participants to report GraphML verification-result witnesses [3, 4] in tandem with each result, and these witnesses must be successfully validated by at least one participating witness validator. Natively, `INFER` does not support the generation of GraphML witnesses. To mitigate this, `INFER-SV` creates generic witnesses: When reporting a violation, it generates a violation witness [4] that represents all possible program paths. When reporting a program safe, it generates a correctness witness [3] that only contains the trivial invariant 'true'. These witnesses do not helpfully guide towards a violation or proof, but are valid according to the SV-COMP rules.

**Participation.** `INFER-SV` participates hors concours in the categories `ReachSafety`, `ConcurrencySafety`, `NoOverflows`, and `SoftwareSystems`. Because of missing support, we exclude `INFER-SV` from categories aimed at float handling, as well as category `MemSafety-MemCleanup`.

<sup>2</sup> <https://fbinfer.com/docs/checker-bufferoverflow>

<sup>3</sup> <https://fbinfer.com/docs/checker-annotation-reachability>

<sup>4</sup> <https://github.com/facebook/infer/issues/763>

<sup>5</sup> <https://fbinfer.com/docs/checker-biabduction>

<sup>6</sup> <https://fbinfer.com/docs/checker-pulse>




Fig. 1: Comparison of the run time (in CPU time seconds) of three SV-COMP '22 medalists and INFER, across all tasks correctly solved by the respective pair

```

1  int main() {
2    if (0) {
3      int x = 0x7fffffff + 1;
4    }
5  }

```

(a) INFER correctly reports safety

```

1  void reach_error() {
2    int x = 0x7fffffff + 1;
3  }
4  int main() {
5    if (0) {
6      reach_error();
7    }
8  }

```

(b) INFER incorrectly reports an alarm

```

1  int main() {
2    int x = 0x7fffffff;
3    int y = -1;
4    while (x > 0) {
5      x = x - 2*y;
6    }
7  }

```

(c) INFER correctly reports an alarm

```

1  int main() {
2    int x = 0x7fffffff;
3    int y = -1;
4    while (x > 0) {
5      x = x - 2*y;
6      y = y + 2;
7    }
8  }

```

(d) INFER incorrectly reports safety

Fig. 2: Examples of INFER's inconsistent results

## 2.2 Strengths of Infer

INFER scales well [6]. This shows in the SV-COMP results: For 6 000 out of 8 000 tasks with a verification verdict, INFER finishes the analysis in less than one second of CPU time. The remaining 2 000 tasks each take less than 100s of CPU time. This means that INFER stays significantly below the time limit of 900s per task. Figure 1 compares the run time of INFER (in CPU-time seconds) to the best SV-COMP '22 tools in the categories that INFER participated in: CPAchecker [11], SYMBIOTIC [8], and VERIABS [12]. Each plot shows the run time for all tasks that are correctly solved by both INFER and the respective other verifier (independent of result validation). It is visible that INFER (y-axis) is significantly faster than the other tools (x-axis) for almost all tasks. This speed makes INFER integrate well in continuous-integration development systems [13, 15].

### 2.3 Weaknesses of Infer

INFER demonstrates low analysis precision. Figures 2a and 2b illustrate a low precision across function calls (intraprocedural analysis): Both programs contain an unreachable, signed integer overflow. The only difference is the indirection in Fig. 2b due to the additional function call. INFER correctly reports Fig. 2a safe, but incorrectly reports an alarm for Fig. 2b. We assume that the intraprocedural analysis of INFER does not check whether `reach_error` is reachable from the program entry. INFER-SV mitigates this issue for property *unreach-call* through the mentioned program transformation, but this imprecision still leads INFER to report wrong alarms across all program properties.

INFER can also show imprecision within a single function. Consider Figs. 2c and 2d: The only change between Fig. 2c and Fig. 2d is the addition of a statement in line 6, `y = y + 2`. This has no influence on the integer overflow in line 5, so both programs contain an overflow. INFER correctly reports the overflow for Fig. 2c, but wrongly reports Fig. 2d safe.

These imprecisions strongly reflect in the SV-COMP results of INFER, leading to many incorrect proofs and alarms.

## 3 Usage

INFER-SV requires Python 3.6 or later. Script `setup.sh` downloads and extracts version 1.1.0 of INFER. From the tool's directory, INFER-SV can be run with the following command:

```
./infer-wrapper.py \
  --data-model {ILP32 or LP64} \
  --property path/to/property.prp \
  --program path/to/program.c \
```

Setting the data model is optional. INFER-SV will print the recognized property and the command line it uses to call INFER. INFER-SV prints the full output of INFER, including all warnings, and the final verification verdict on the last line. The verification verdict can be `true`, `false`, `unknown` or `error`.

## 4 Conclusion

The participation of INFER in SV-COMP allows an objective comparison with other verifiers for C. This shows that the selected analyses of INFER are very efficient, but suffer from strong imprecision on the considered benchmark tasks.

**Contributors.** INFER<sup>7</sup> is developed by Facebook and the open-source community under the MIT license, and INFER-SV<sup>8</sup> is developed under the Apache 2.0 license at the Software and Computational Systems Lab at LMU Munich, led by Dirk Beyer.

<sup>7</sup> <https://github.com/facebook/infer>

<sup>8</sup> <https://gitlab.com/sosy-lab/software/infer-sv>

**Funding Statement.** This work was funded in part by the Deutsche Forschungsgemeinschaft (DFG) – 418257054 (Coop).

**Data Availability Statement.** All data of SV-COMP 2022 are archived as described in the competition report [1] and available on the [competition web site](#). This includes the verification tasks, results, witnesses, scripts, and instructions for reproduction. The version of our verifier as used in the competition is archived together with other participating tools [2].

## References

1. Beyer, D.: Progress on software verification: SV-COMP 2022. In: Proc. TACAS. Springer (2022)
2. Beyer, D.: Verifiers and validators of the 11th Intl. Competition on Software Verification (SV-COMP 2022). Zenodo (2022). <https://doi.org/10.5281/zenodo.5959149>
3. Beyer, D., Dangl, M., Dietsch, D., Heizmann, M.: Correctness witnesses: Exchanging verification results between verifiers. In: Proc. FSE. pp. 326–337. ACM (2016). <https://doi.org/10.1145/2950290.2950351>
4. Beyer, D., Dangl, M., Dietsch, D., Heizmann, M., Stahlbauer, A.: Witness validation and stepwise testification across software verifiers. In: Proc. FSE. pp. 721–733. ACM (2015). <https://doi.org/10.1145/2786805.2786867>
5. Beyer, D., Spiessl, M.: The static analyzer FRAMA-C in SV-COMP (competition contribution). In: Proc. TACAS (2). Springer (2022)
6. Calcagno, C., Distefano, D., Dubreil, J., Gabi, D., Hooimeijer, P., Luca, M., O’Hearn, P.W., Papakonstantinou, I., Purbrick, J., Rodriguez, D.: Moving fast with software verification. In: Proc. NFM. pp. 3–11. LNCS 9058, Springer (2015). [https://doi.org/10.1007/978-3-319-17524-9\\_1](https://doi.org/10.1007/978-3-319-17524-9_1)
7. Calcagno, C., Distefano, D., O’Hearn, P.W., Yang, H.: Compositional shape analysis by means of bi-abduction. ACM **58**(6), 26:1–26:66 (2011). <https://doi.org/10.1145/2049697.2049700>
8. Chalupa, M., Rečtáček, A., Mihalkovič, V., Zaoral, L., Strejček, J.: SYMBIOTIC 9: Parallelism and invariants (competition contribution). In: Proc. TACAS (2). Springer (2022)
9. Chowdhury, A.B., Medicherla, R.K., Venkatesh, R.: Verifuzz: Program aware fuzzing - (competition contribution). In: Proc. TACAS, part 3. pp. 244–249. LNCS 11429, Springer (2019). [https://doi.org/10.1007/978-3-030-17502-3\\_22](https://doi.org/10.1007/978-3-030-17502-3_22)
10. Clarke, E.M., Kröning, D., Lerda, F.: A tool for checking ANSL-C programs. In: Proc. TACAS. pp. 168–176. LNCS 2988, Springer (2004). [https://doi.org/10.1007/978-3-540-24730-2\\_15](https://doi.org/10.1007/978-3-540-24730-2_15)
11. Dangl, M., Löwe, S., Wendler, P.: CPACHECKER with support for recursive programs and floating-point arithmetic (competition contribution). In: Proc. TACAS. pp. 423–425. LNCS 9035, Springer (2015). [https://doi.org/10.1007/978-3-662-46681-0\\_34](https://doi.org/10.1007/978-3-662-46681-0_34)
12. Darke, P., Agrawal, S., Venkatesh, R.: VERIABS: A tool for scalable verification by abstraction (competition contribution). In: Proc. TACAS (2). pp. 458–462. LNCS 12652, Springer (2021). [https://doi.org/10.1007/978-3-030-72013-1\\_32](https://doi.org/10.1007/978-3-030-72013-1_32)
13. Distefano, D., Fähndrich, M., Logozzo, F., O’Hearn, P.W.: Scaling static analyses at Facebook. Commun. ACM **62**(8), 62–70 (2019). <https://doi.org/10.1145/3338112>

14. Distefano, D., O'Hearn, P.W., Yang, H.: A local shape analysis based on separation logic. In: Proc. TACAS. LNCS, vol. 3920, pp. 287–302. Springer (2006). [https://doi.org/10.1007/11691372\\_19](https://doi.org/10.1007/11691372_19)
15. Harman, M., O'Hearn, P.W.: From start-ups to scale-ups: Opportunities and open problems for static and dynamic program analysis. In: Proc. SCAM. pp. 1–23. IEEE (2018). <https://doi.org/10.1109/SCAM.2018.00009>
16. Malík, V., Schrammel, P., Vojnar, T.: 2LS: Heap analysis and memory safety (competition contribution). In: Proc. TACAS (2). pp. 368–372. LNCS 12079, Springer (2020). [https://doi.org/10.1007/978-3-030-45237-7\\_22](https://doi.org/10.1007/978-3-030-45237-7_22)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

