

## **ETAPS Poster Book**

### **Collection of posters presented at ETAPS 2025**

Hamilton, May 3-8, 2025

FASE 2025 Posters

## **ByteBack: Deductive Verification** at the Level of JVM Bytecode

#### Marco Paganoni Carlo A. Furia •

Software Institute, USI Lugano, Switzerland

### **Deductive Verification**



### **Supported Java Features**

|                             | Java    |  | Suppo   | rt   |
|-----------------------------|---------|--|---|--|
| Feature                     | Version | KeY  | OpenJML   | ByteBack   |
| Generic classes             | 5       |  | <b>~</b>  | <b>~</b>   |
| Enhanced for loop           | 5       | <ul> <li>Image: A second s</li></ul> | <ul> <li>Image: A set of the set of the</li></ul> | <ul> <li>Image: A second s</li></ul> |
| Variable arguments          | 5       | <ul> <li>Image: A second s</li></ul> | ×   | <ul> <li>Image: A second s</li></ul> |
| Generic type inference      | 7       | ×  | ×   | <ul> <li>Image: A second s</li></ul> |
| Try-with-resources block    | 7       | ×  | ×   | <ul> <li>Image: A second s</li></ul> |
| Multi-catch block           | 7       | <ul> <li>Image: A second s</li></ul> | ×   | <ul> <li>Image: A set of the set of the</li></ul>  |
| Default methods             | 8       | <ul> <li>Image: A second s</li></ul> | <ul> <li>Image: A set of the set of the</li></ul> | <ul> <li>Image: A set of the set of the</li></ul>  |
| Local type inference        | 10      | ×  | ×   | <ul> <li>Image: A second s</li></ul> |
| switch expressions          | 12      | ×  | ×   | <ul> <li>Image: A second s</li></ul> |
| yield in switch expressions | 13      | ×  | ×   | <ul> <li>Image: A set of the set of the</li></ul>  |
| instanceof pattern matching | 14      | ×  | ×   | <b>~</b>   |
|                             |         |  |   |  |

### ByteBack's Architecture



https://github.com/atom-sw/byteback

### **BBlib: Specification in ByteBack**

### Experiments with ByteBack

| <pre>@Require("no_twos") // @require(forall j : int • xs[j] ≠ 2)<br/>@Ensure("nonnegative") // @ensure(return ≥ 0)<br/>public int onesMinusTwos(int xs) {<br/>var r = 0;<br/>for (var x : xs) {<br/>invariant(gte(r, 0)); // @invariant(r ≥ 0)<br/>r += switch (x) {<br/>case 1 -&gt; 1;</pre> | 120 Verified Programs |  |                                |                                  |                      |  |  |
|--|-----------------------|--|--------------------------------|----------------------------------|----------------------|--|--|
|  |                       | Langua                                   | age Coun                       | t Size (LO                       | C)                   |  |  |
| <pre>var r = 0;<br/>for (var x: xs) {<br/>invariant(gte(r, 0)); // @invariant(r ≥ 0)<br/>r += switch (x) {<br/>case 1 -&gt; 1;<br/>case 2 -&gt; -1;<br/>default -&gt; 0;</pre>   |                       | Java 8<br>Java 17<br>Scala 2<br>Kotlin 1 | 58<br>7 23<br>2.3 23<br>1.8 18 | 8 817<br>3 648<br>1 222<br>8 230 | 71<br>39<br>21<br>03 |  |  |
| ),<br>}<br>return r;   | Benchmarks            |  |                                |                                  |                      |  |  |
| }  |                       | Size                                     | (LOC)                          | Tir                              | me ( <i>s</i> )      |  |  |
| <pre>@Predicate public boolean no_twos(int[] xs)</pre>   |                       | Source                                   | Boogie                         | Encoding                         | Verification         |  |  |
| <pre>{ int j = Binding.integer(); return forall(j, neq(xs[j], 2)); }</pre>   | Total                 | 1984                                     | 1031874                        | 66.3                             | 104.7                |  |  |
| <pre>@Predicate public boolean nonnegative(int[] xs, int r) {     return is ( 0); }</pre>  | Average               | 159                                      | 8598                           | 552.7                            | 872.5                |  |  |

an nonnegative(int[] xs, int r) { return gte(r, 0); }



### Human-Robot Interaction in the Design and Verification of Robotic Systems



The Pilot requests the battery status of the drone from the Handheld. The Handheld returns it to the Pilot who then, after the passage of nondeterministic time, decides if the status is ok.

· Development of a tool to automatically translate from RoboScene into a mathematical notation (CSP) · Connection of the human, software and hardware

models

UNIVERSITY of Vork THALES

- · Receive counterexamples in the case of property failures (FDR)

Properties of the SAR example are derived from human factors models

Royal Academy of Énaineerina



Engineering and Physical Sciences **Research Council** 

### Hybridize Functions: A Tool for Automatically Refactoring Imperative Deep Learning **Programs to Graph Execution**

Raffi Khatchadourian<sup>1,2</sup> Tatiana Castro Vélez<sup>2</sup>

Mehdi Bagherzadeh<sup>3</sup>

<sup>1</sup>CUNY Hunter College, USA (ponder@hunter.cuny.edu) <sup>2</sup>CUNY Graduate Center, USA <sup>3</sup>Oakland University, USA

Anita Raja<sup>1,2</sup>

### Introduction

- As Deep Learning (DL) datasets grow, efficiency becomes essential to support responsiveness [16].
- Traditionally, DL frameworks embraced deferred execution-style DL code for fast execution.
- Hybrid approaches [2, 8, 13] execute imperative DL programs quickly.

### Hybridization





In TensorFlow [1], AutoGraph [13] can enhance run-time performance by decorating (annotating) appropriate Python function(s) with @tf.function (Fig. 1).

### Problems with Hybrid Approaches

- Require non-trivial metadata [12].
- Exhibit limitations and known issues with native program constructs [9].
- Are difficult to use correctly and efficiently (e.g., avoiding side-effects) [4].
- Developers manually specifying which functions are converted.

### Insight

Although imperative DL code typically executes sequentially, hybridization resembles parallelizing traditional sequential code.

### Automated Tool

We design and implement a fully automated, open-source refactoring tool named HYBRIDIZE FUNCTIONS [11] that transforms otherwise eagerly-executed imperative (Python) DL code for enhanced performance.

### Contributions

- Refactoring approach for automatically converting imperative DL code to graphs.
- Novel tensor analysis for imperative DL.
- Fully automated, open-source tool implemented as a PyDev [15] Eclipse [7] IDE plug-in that integrates static analyses from WALA [14] and Ariadne [6].



### Figure: Overall architecture

 Eclipse is leveraged for its existing, well documented and integrated refactoring framework and test engine [3], including transformation APIs (e.g., ASTRewrite), refactoring preview pane (Fig. 1), precondition checking (e.g., Refactoring.

checkInitialConditions(),

Refactoring.

checkFinalPreconditions()), and refactoring testing (e.g., RefactoringTest).

- PyDev used for efficient program entity indexing, extensive refactoring support [3], and that it is completely open-source for all Python development.
- WALA is used for static analyses, such as ModRef, for which we built our side-effect analysis upon.
- Ariadne, which depends on WALA, is used for its Python and tensor analysis, including type inference and (*TensorFlow*) library modeling.

### Challenges Addressed

- Reworked much of the existing Java (JDT) refactoring tooling to work with Python.
- Integrated Ariadne with PyDev due to its excellent and long-lived refactoring support for Python, including refactoring preview pane, element GUI selection, and refactoring undo history.
- Augmented Ariadne to analyze imperative Deep Learning (Python) code by vastly expanding the XML summaries to support a wide variety of popular TensorFlow 2 APIs.
- Added support for Python constructs commonly used in modern imperative DL programs.
- Correlated varying intermediate representations (IRs) with the original Python source code.

- Modernizing Ariadne: New Enhancements
- Python module packages.
- Wild card imports.

Nan Jia<sup>2</sup>

- Intra-package references (relative imports; from .. import X).
- Package initialization scripts.
- Automatic unit test entry points discovery.
- Non-scalar tensor dataset [10] iteration.
- Modeling of additional libraries.
- Static and class methods analysis.
- Analysis of custom decorators.
- Callable object (functor) analysis (used in Keras).

### **Evaluation Summary**

- We applied our approach to 19 open-source Python imperative DL programs of varying size and domain, with thousands of source lines of code ranging from 0.12 to 36.72.
- Our tool considered 766 Python functions, automatically refactoring 42.56% despite being highly conservative.
- During a run-time performance evaluation, we measured an average relative model training speedup of 2.16 (memory consumption measurement pending).
- Differences in model accuracy and loss before and after refactoring were negligible.

### Conclusion

- Open-source, automated refactoring PyDev Eclipse plug-in, HYBRIDIZE FUNCTIONS, assists developers with writing optimal imperative DL Python code.
- Integrates an Eclipse refactoring with WALA Ariadne Python static analyses.

### Future Work

- Explore incorporating advanced container-based analyses.
- Automatically split functions.

### References

- Abadi, M. et al.: TensorFlow: A System for Large-Scale Machine Learning. In: OSDI (2016) Abache, Hybridize. Apache MXNet documentation. (2021). https://mxnet. //mi/umithom/docs/lutorials/packages/gluon/blocks/hybridize.ht
- ating refactoring support into a Java development tool Bäumer, D. et al.: "Integ Castro Vélez, T. et al.: Challenges in Migrating Imperative Deep Learning Programs to Graph Execution: An Empirical Study. In: MSR. MSR '22. ACM (2022). https://doi.org/10.1145/3804840\_3809846
- Chollet, F.: Deep Learning with Python. Dolby, J. et al.: Ariadne. Analysis for Machine Learning Programs. In: MAPL, pp. 1-10. ACM (2018)
- Eclipse Foundat Eclipse Foundation, Eclipse IDE. (2024). ht Facebook Inc., PyTorch. TorchScript. en. ( //eclipseide.org/ (visited on 09/10/2024)
  . https://pvtorch.org/docs/stable/jit.html
- Google LLC, Better performance with tf.function Google LLC, tf.data.Dataset. TensorFlow. Vers
- 11. Hybridize-Functions-Refactoring. (2
- Jeong, E. et al.: Speculative Symbolic Graph Execution of Imperative Deep Learning Prog Oper. Syst. Rev. 53(1), 26-33 (2019). https://doi.org/10.1145/3352020.3352025 13. Moldovan, D. et al.: AutoGraph: Imperative-style Coding with Graph-based Performance. (2019). arX
- T.J. Watson Libraries for Analysis. (2024). https://github.com/wala/WALA (visited on 09/10/2024 critical data: 2012.04.05T19:57:027
- 15. Zadrozny, F.: PyDev. (2023). https F.: PyDev. (2023). https://www.pydev.org (visited on 05/31/2023) et al.: HARP: Holistic Analysis for Refactoring Python-Based Analytics Programs. In: ICSE (2020)

Acknowledgments This material is supported in part by the National Science Foundation under awards CCF 2200343, CNS 2213763, and CCF 2343750.

International Conference on Fundamental Approaches to Software Engineering, May 3-8, 2025, Hamilton, Canada

### SYMBOLIC STATE PARTITIONING FOR REINFORCEMENT LEARNING Mohsen Ghaffari<sup>1</sup>, Mahsa Varshosaz<sup>1</sup>, Einar Broch Johnsen<sup>2</sup>, Andrzej Wąsowski<sup>1</sup> <sup>1</sup> IT-University of Copenhagen, Copenhagen, Denmark <sup>2</sup> University of Oslo, Norway Talk: FASE, Tuesday, 11:00

Full Paper

### Problem: How to partition the state space for training an agent efficiently using Tabular Reinforcement Learning?





**Capturing nonlinear dependencies** between state components, and **finding narrow parts**.

### SymPar

- •The environment is modeled as a computer program that implements a singlestep transition function, producing the next state and the corresponding reward,
- For each concrete action, the environment simulator is symbolically executed to generate a set of partitions over the state space. SymPar then computes the intersection of these partitions as its final output,
- **Symbolic Execution** extends normal execution by running the operators of a language using symbolic variables and producing symbolic formulas (PC) as output,
- Symbolic execution computes semantic based partition.

### **Experiments Results**

### How does granularity of partition affect learning performance?

Increasing the search depth of the symbolic executor leads to finergrained partitions produced by SymPar. This higher granularity enhances learning performance, as reflected in improved accumulated rewards.

Simple

UNIVERSITY OF OSLO

Maze

### How does SymPar scale with state space sizes?

The number of partitions generated by Sym-Par remains constant as the state space grows, demonstrating scalability with respect to state space size.





IT UNIVERSITY OF COPENHAGEN

FoSSaCS 2025 Posters

# **Fair Quantitative Games**

Ashwani Anand, Satya Prakash Nayak, Ritam Raha, Irmak Sağlam, and Anne-Kathrin Schmuck

Can we guarantee the robot will put

a block at the mark, and remove it, infinitely often?

Oualitative

Safety, Parity, Rabin...

Fairness is easy on qualitative games

Can be solved in same time

complexity as the original game [1, 2]





Environment

### Contributions

|               | Determined? | Complexity<br>(Pseudopolynomial) | Reduction   |
|---------------|-------------|----------------------------------|---|
| 1-fair MP     | Yes         | $O(n^3mW)$                       | To MP on 6n nodes and max absolute weight   |
| 2-fair MP     | Yes         | $O(n^3mW)$                       | To MP on 6n nodes and max absolute weight   |
| 1-fair Energy | Yes         | $O(n^4mW)$                       | To Energy on 8n nodes and max absolute weight   |
| 2-fair Energy | No          | 0(n <sup>3</sup> mW)             | Player 1 winning region reduces to that of an energy on the same<br>graph, Player 2 winning region reduces to that of 2-fair MP game<br>on the same graph |

### Gadgets for Reducing Fair Games to Standard games

For each fair node *v* in the 1-fair MP game

replace v with the following v-gadget



### **Fairness in Synthesis**



Can we guarantee the robot will put a block at the mark, and remove it, infinitely often, without the battery running out, for some initial value of the battery?

> Ouantitative Energy, Mean-payoff

#### What about fairness on quantitative games?

Fairness on system can be solved in **super-exponential time** using current approaches, whereas there is no known approach for fairness on environment.

### **Fairness in Quantitative Games**

A play  $\rho$  is **fair** iff for every node  $v \in inf(\rho)$  that has fair (dashed) outgoing edges  $E^{f}(v) \neq \emptyset$ ,  $E^{f}(v) \subseteq inf(\rho)$ .

1-Fair Games: Player 1 nodes have fair outgoing edges

2-Fair Games: Player 2 nodes have fair outgoing edges

#### 1-Fair Mean Payoff

### 1-Fair Energy

Does there exist a strategy  $\sigma$  such that, long run average payoff of every  $\sigma$ -play is non-negative AND the play is fair? Does there exist an initial credit  ${\bf c}$  and a strategy  $\sigma$  such that, total energy level along every  $\sigma$ -play stays non-negative AND the play is fair?

#### 2-Fair Mean Payoff

Does there exist a strategy  $\sigma$  such that, long run average payoff of every  $\sigma$ -play is non-negative OR the play is NOT fair? Does there exist an initial credit **c** and a strategy  $\sigma$  such that, total energy level along a play stays **non-negative OR the play is NOT** fair?

2-Fair Energy

### 2-fair energy games are not determined



A node is won by Player 1 if there exists a Player 1 strategy  $\sigma$  and a credit c s.t. every  $\sigma$  – play is won by Player 1 for credit c.

A node is won by Player 2 if there exists a Player 2 strategy  $\pi$  s.t. every  $\pi$  – play's total payoff goes below –c for every credit c.

[1] Banerjee, T., Majumdar, R., Mallik, K., Schmuck, A., Soudjani, S.: Fast symbolic algorithms for omega-regular games under strong transition fairness. TACAS'22 [2] Hausmann, D., Piterman, N., Saglam, I., Schmuck, A.: Fair ømega-regular games. FoSSaCS'24



### MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS

**TACAS 2025 Posters** 

# Fixed Point Certificates for Reachability and Expected Rewards in MDPs

Krishnendu Chatterjee, Maximilian Weininger, Tim Quatmann, Tobias Winkler, Maximilian Schäffeler, and Daniel Zilken







- cates as a new standard for Certified Model Checking
- We compute certificates with rational value vectors and check them with exact, arbitrary precision arithmetic
- Scalable to  $\approx 10^6$  states with a runtime within  $\approx 30$  sec.
- Soundness formalized in Isabelle/HOL
- CT: daniel.zilken@cs.rwth-aachen.de



Technische Universität München





## GPU accelerated probabilistic model checking

Jan Heemstra (j.h.heemstra@tue.nl), Anton Wijs (a.j.wijs@tue.nl)

### Objectives

- $\bullet$  Perform probabilistic model checking **entirely** on a  $\mathbf{GPU}.$
- Previous achievements with the GPU explore model checker:
  - -Use hardware acceleration to speed process up [4]
  - Verify specification adherence with LTL formulae [3].
- New: compute the **probability** of invalid system behavior with PCTL formulae [1, 2].

### Matrix construction

- Goal: construct (sparse) transition matrix on a GPU.
- Input: explored states are chaotically distributed over hash table.
- Sparse matrix is stored in memory using **CSR format**.



- $\bullet$  Memory usage is main constraint.
- On-chip (a.k.a. shared) memory is fast and low latency.



• Two prefix sums are needed: for row offsets and column indices.

• Interweaved prefix sum optimizing both memory access latency and usage.



- After matrix construction, the CUSPARSE library is used to perform matrix-vector multiplications, for the verification of PCTL formulae.
- $\bullet$  Up to  $855\times$  faster than Storm, fastest on larger models.



TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY





### References

- Heemstra, J., Osama, M., Wijs, A.: Towards End-to-End GPU Acceleration of PCTL Model Checking. Springer, Cham (2025)
- [2] Heemstra, J., Wijs, A.: GPUexplore-prob: Markov Chain State Space Construction and Verification with GPUs. In: TACAS (2025)
- [3] Osama, M., Wijs, A.J.: Hitching a Ride to a Lasso: Massively Parallel On-The-Fly LTL Model Checking. In: TACAS. Springer (2024)
- [4] Wijs, A.J., Osama, M.: GPUexplore 3.0: GPU Accelerated State Space Exploration for Concurrent Systems with Data. In: SPIN. Springer (2023)

#### https://github.com/egolf-cs/tacas25

### ACCELERATING PROTOCOL SYNTHESIS AND DETECTING UNREALIZABILITY WITH INTERPRETATION REDUCTION

Derek Egolf, Stavros Tripakis

Northeastern University, Boston



### **Key Contributions**

- Synthesize symbolic distributed protocols represented in TLA<sup>+</sup> [Lamport].
- $\bullet\, {\rm Improve \ state \ of \ art \ in \ TLA^+}$  synthesis (100x).
- $\bullet$  Synthesize a lock protocol "from scratch."
- $\bullet$  Halt when no solution: unrealizability.
- $\bullet$  New search space reduction technique: Interpretation Reduction.
- Improved counterexample generalization for pruning.

### Sketching [Solar-Lezama]

Given an incomplete sketch with "holes," find a correct completion.

| Example sketch:            | Example completion                                   |
|----------------------------|--|
| Send(src, dst) :=          | Send(src, dst) :=                                    |
| $\land ???_1$              | $\land$ has_lock[src]                                |
| $\land message' = ???_2$   | $\land message' = message \cup \{(src, dst)\}$       |
| $\land has\_lock' = ???_3$ | $\land has\_lock' = has\_lock[src \leftarrow false]$ |
| Receive(src, dst) :=       | Receive(src, dst) :=                                 |
| $\land ???_4$              | $\land$ (src, dst) $\in$ message                     |
| $\land message' = ???_5$   | $\land message' = message \setminus \{(src, dst)\}$  |
| $\land has\_lock' = ???_6$ | $\land has\_lock' = has\_lock[dst \leftarrow true]$  |
|                            |  |



### **Experimental Results**



n = 171; worse: 11; scythe TO: 47; poly TO 15

**Unrealizable Experiments** (n = 123)

• **Polysemist** (new): 
$$TO = 43$$
 / HALT = 80

- Usually halted in 
$$< 60$$
 seconds

-Did not TO unless Scythe did

### Counterexample-Guided Inductive Synthesis (CEGIS)

Our approach uses standard CEGIS technique [Solar-Lezama].



### Key Technical Ideas for Learner

- Naive learner: ignore cex, enumerate all protocols many expressions; model checking expensive
- Pruning constraints: generalize counterexamples discard protocols before model checking
- Equivalence reduction: do not use equiv. sub-expressions avoid enumerating protocols in the first place
- $\bullet$  Interpretation reduction: coarse, dynamic equivalence relation

### Counterexample Generalization

 $\begin{array}{l} \textbf{safety cex: } [x \mapsto a] \xrightarrow{A_1} [x \mapsto b] \\ \textbf{Sketch: } (A_1 := ???_1 \land x' = ???_2) \\ \textbf{Bad Completion: } (A_1 := \textit{true} \land x' = b) \\ \textbf{Bad Completion: } (A_1 := x = a \land x' = b) \\ \textbf{Good Completion: } (A_1 := x = a \land x' = a) \\ \textbf{Good Completion: } (A_1 := x \neq a \land x' = b) \\ \textbf{Good Completion: } (A_1 := x \neq a \land x' = b) \end{array}$ 

### Pruning Constraint:

 $\tau$ 

$$T_{cex} := ???_1([x \mapsto a]) \neq true \lor ???_2([x \mapsto a]) \neq b$$

- In general, many completions violate  $\pi_{cex}$
- Checking  $P \vDash \pi_{cex}$  is much cheaper than model checking.
- We generalize deadlock, safety, and liveness violations.
- Prior work [FMCAD'24] uses less exact pruning constraints for dead-lock/liveness.

### Interpretation Reduction

Absolute Equivalence, e.g:  $x + y \equiv y + x$ Interpretation Equivalence, e.g:  $x + y \equiv_{\mathcal{A}} x + x$ , where  $\mathcal{A} = \{[x \mapsto 0, y \mapsto 0]\}$ 

 $\bullet \mathcal{A}$  comes from pruning constraints. E.g.,

$$\Pi = \{ ???_1([x \mapsto 0, y \mapsto 1]) \neq 1 \} \rightarrow \mathcal{A} = \{ [x \mapsto 0, y \mapsto 1] \}$$

- $\bullet$  Suppose we've enumerated y; enumerate x+y? No:  $y\equiv_{\mathcal{A}}x+y.$
- Avoids enumerating all "super-expressions" of x + y, e.g., x + x + y.
- $\bullet$  Coarse eq. relation makes detecting unrealizability faster.
- [FMCAD'24] uses absolute equivalence for reduction.

**Theorem**: If  $e_1 \equiv_{\mathcal{A}} e_2$  and  $e_1$  enumerated, skipping  $e_2$  does not compromise completeness

References Lamport, L.: Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers. Addison-Wesley (2002). Solat-Learna, A.: Program Sketching, Int. J. Softw. Tools Technol. Transf. (2013). Egolf, D., Tripaks, S.: Efficient Synthesis of Distributed Protocols by Sketching, FMCAD (2024).

## ETH zürich

## Pushing the Limit: Verified Performance-Optimal Causally-Consistent Database Transactions





Dr. Christoph Sprenger



Dr Si Liu



Luca Multazzu



Prof. Dr. David Basin

### 1 Motivation

- Distributed databases (key-value stores)
- **Isolation** of concurrent transactions, realized by concurrency control protocols.
- Spectrum of isolation levels:



- Trade-off: isolation vs performance
- TCC previously conjectured to be the strongest achievable isolation level for performance-optimal reads in the presence of transactional writes.
- We refute the conjecture and push the limit to TCCv with our novel protocol Eiger-PORT+.
- Concurrency control protocols are highly complex and prone to design errors and isolation bugs. → Deductive verification

### Contributions

- Eiger-PORT+
  - > Stronger isolation guarantee
  - > Superior performance
- \* Protocol verification in Isabelle/HOL
  - Refinement
  - ➤ Reduction

### 2 Abstract Model: Isolation Level

We specify isolation as an abstract model parameterized by an Isolation Level (IL).









- read-only transactions
- → own write above gst
- → or last write below gst
- write-only transactions: two-phase commit



### 4 Correctness Proof & Invariants

#### Proof guarantee:

refinement (reach (protocol))  $\subseteq$  reach (IL)

- refinement mapping:
  - r : K and U reconstructed as shown above *π*: client\_write\_commit and client\_read\_done mapped to *Atomic commit*
- proof obligations:
- canCommit<sub>IL</sub>: needs invariants (below) LWW: needs reduction



• The relation of different timestamps in the model and finding invariants





### **5 Inverted Commits & Reduction**

Inverted commits: pairs of client commits in protocol executions not ordered by **commit timestamps**.



Can occur for causally independent concurrent transactions.

**Problem:** Inverted commits would require *inserting* rather than *appending* a transaction's version to the version list.

Can not be simulated by the abstract model.

Hence, refinement alone is not enough for verifying the protocol.

#### Solution:

• We introduce a restricted protocol model that doesn't produce inverted commits.



- We use reduction to transform any protocol execution into one of the restricted model such that: reach (protocol) = reach (restricted protocol).
- This is achieved by commuting independent concurrent events to eliminate inverted commits. (see arrows on the execution above)

### **6** Conclusions and Discussion

- Our Eiger-PORT+ protocol provides
   TCCv, thus refuting an open conjecture.
   Eigen DORT+ subsetformer state of the set



- Refinement is not always enough
- We deductively verify that Eiger-PORT+ satisfies TCCv, using a combination of refinement and reduction.



Previous Publications & Open Sourced Tools

Y.-H. Tsai, J.-H. R. Jiang, and C.-S. Jhang, "Bit-slicing the Hilbert space: Scaling up accurate quantum circuit simulation," in *Proc. DAC*, 2021, pp. 439–444.
 C.-Y. Wei, Y.-H. Tsai, C. -S. Jhang, and J.-H. R. Jiang, "Accurate BDD-based Unitary Operator Manipulation for Scalable and Robust Quantum Circuit Verification," in *Proc. DAC*, 2021, pp. 439–444.

- 2022, pp. 523-528.
- 3. T.-F. Chen, J.-H. R. Jiang, and M.-H. Hsieh, "Partial Equivalence Checking of Quantum Circuits," in *Proc. QCE*, 2022.

a) Simulator – SliQSim https://github.com/NTU-ALComLab/SliQSim b) Equivalence Checker – SliQEC https://github.com/NTU-ALComLab/SliQEC

## Synthesis with Guided Environments

Orna Kupferman and Ofer Leshkowitz

Hebrew University of Jerusalem, Jerusalem, Israel

### Synthesis of Reactive Systems

I = Finite set of input si O = Finite set of output

#### Reactive System -

- Environment generates assignments to input signals  $i_1, i_2, i_3, \dots \in 2^{1}$
- System generates assignments to output signals  $o_1, o_2, o_3, ... \in 2^0$
- Outputs generated online:  $o_k$  is determined by  $i_1, i_2, ..., i_k$ .
- Synthesis - $\varphi = LTL$  formula
- Given a specification  $\varphi$  of "good computations", automatically construct a reactive system that realizes  $\varphi$ .

#### Dealing with partial visibility (Back to Reactive Systems...)

When visiting your doctor do you -







### Dealing with partial visibility

- The assignment to  $H \subseteq I$  is hidden from the system.
- System should satisfy  $\varphi$  for all assignments to H.
- Simple specification become unrealizable.



 $\varphi = i \leftrightarrow o$   $I = \{i\}, O = \{o\} \text{ and } H = \{i\}.$ 

Example:

#### Transducer with a Guided Environment

### A TGE sends **"programs"** used by Environment to assign guided outputs.



#### Environment needs memory.

#### Memory Monotonicity

- More visibility \ guidance ∖ Less memory. 1. Increasing visibility (viewing also  $h \in H$ ) can only reduce needed memory. Assignment to h can be hard-coded into programs
- 2. Increasing guidance (guiding also  $c \in C$ ) can only reduce needed memory. Assignment to c can be hard-coded into programs

- What's going on? • Magician -1. Gains partial information 2. Guides fish step-by-step. • Fish -1. Picks p adversarially
  - 2. Follows instructions -Remembers values & Performs calculations.

#### Transducer with a Guided Environment

<u>A round of interaction with a TGE:</u> • Env evaluates p(m, h) = (m', g).  $(l = V \cup H, 0 = C \cup G)$  Updates memory to m'. Sets guided out to g Env generates:

### • Visible input v & Hidden input h.

 TGE reads visible input v: Updates state to s'. Generates -

 $\begin{array}{l} \text{Controlled output } c \text{ \& Program } p. \\ c \in 2^{\mathcal{C}} \text{ and } p \text{:} M \times 2^{\mathcal{H}} \rightarrow M \times 2^{\mathcal{G}}. \end{array}$ 

Think of a two-digit prime nu

What's the first digit of the

Subtract first digit from p<sup>2</sup>. What's the second digit of the result?

p + 2 is prime as well !!

6

#### Solving SGE - Synthesis with Guided Environments

**The problem:** Synthesize a TGE that realizes  $\varphi$  with *H* hidden, *G* guided and at most *k* environment memories or determine that there is no such TGE.

#### **Theorem:** (1) A tree $\mathcal{T}$ is accepted by $A^{\varphi}_{M,H,G}$ iff it realizes $\varphi$ . (2) $A^{\varphi}_{M,H,G}$ is of size $|M| \cdot \exp(|\varphi|)$ . **Reduce** synthesis to the **non-emptiness** of $A^{\varphi}_{M,H,G}$ :

1. Search for a finite witness  ${\mathcal T}$  for non-emptiness of  $A^{\varphi}_{M,H,G}$  in EXPTIME. 2. If empty - Not realizable.

3. Otherwise -  ${\mathcal T}$  describes a TGE realizing  ${\varphi}$  with  ${\cal H}$  hidden and  ${\cal G}$ 

auided.

**Corollary:** SGE is 2EXPTIME in  $|\phi|$  and EXPTIME in |M|.

#### Non-Zero-Sum Games with **Multiple Weighted Objectives**



We introduce and study non-zero-sum multi-player games with weighted multiple objectives.

In these games, the objective of each player consists of a set  $\alpha$  of underlying objectives and a weight function w:  $2^{\alpha} \rightarrow \mathbb{Z}$ that maps each subset X of  $\alpha$  to the utility of the player when exactly all the objectives in X are satisfied. The weight functions lift the setting of non-zero-sum multi-player games to the general quantitative case, allowing a rich reference to the underlying objectives.

Settings: · Several interacting agents (robots, users)

each having an objective
 each objective is composed of weighted sub-objectives



k-player weighted multiple objectives

Set of objectives  $\alpha = \{\alpha_1, \alpha_2, ..., \alpha_m\} \quad \alpha_l \subseteq V \quad (\alpha_l \text{ is a}$ Büchi objective) Set of weight functions  $\{w_1, w_2, ..., w_k\}$ (one for each player in the game)  $w_i: 2^{\alpha} \rightarrow \mathbb{Z}$ 

For a play  $\rho$ :

 $\begin{aligned} -\operatorname{sat}(\rho,\alpha) &\subseteq \alpha \\ -\operatorname{For} \operatorname{every} i \in [k], \operatorname{value}(\rho,\alpha,w_i) = w_i(\operatorname{sat}(\rho,\alpha)) \end{aligned}$ The utility of Player i

Stable profiles in weighted k-player games: Zero-sum games: No overlap among objectives, find the winning players

Non-zero-sum games: Overlap among objectives, find stable profiles

A predicate describes desired utilities of the play. For example:  $u_1 > 2 \land (u_2 < 3 \lor u_3 \ge 1)$ 

For a set  $S \subseteq [k]$  of system players, and a predicate P. d NE with Desired Utilities (DNE) problem is to The Partially-Fix return an S-fixed NE  $\pi$  that satisfies P.

### There is a warehouse, P1 and P2 are robots who wants to access the warehouse to move inventory. P3 is a guard making sure all is in order.



Take a picture to download the full paper

#### Our contribution:

| Limitation on the weights   | Zero-sum games          | Non-zero-sum DNE                  |  |  |
|-----------------------------|-------------------------|-----------------------------------|--|--|
| No limitation               | PSPACE-complete [KS 24] | PSPACE-complete (already for k=2) |  |  |
| Non-decreasing              | co-NP-complete [KS 24]  | NP-complete (already for k=2)     |  |  |
| Additive                    | PSPACE-complete         | PSPACE-complete (already for k=2) |  |  |
| Almost-positive<br>additive | co-NP-complete          | NP-complete (already for k=2)     |  |  |

#### Applications of k-player weighted multiple objectives

We study the extension of the game by payments, with which players can incentivize each other to follow strategies that are beneficial for the paying player. We show how such payments can be used in order to repair systems.

#### Games with Payments

 $\{\alpha_1, \alpha_2, \dots, \alpha_k\} \quad \alpha_i \subseteq V$ Set of Buchi objectives, one for each player

 $\{\mathsf{R}_1,\,\mathsf{R}_2,...,\,\mathsf{R}_k\} \quad \mathsf{R}_i \in \,\mathbb{N}$  $\mathbf{R_i}$  the reward Player i receives when objective  $\alpha_i$  is satisfied

A payment function  $p\colon [k]\times [k]\to\mathbb{N}$  p(i,j) the amount Player i commits to pay Player j when  $\alpha_i$  is satisfied

#### For a play ho, W( ho) is the set of objectives satisfied by ho



The monetary-based system-repair problem is given a set of system players S, and a predicate P to decide if there exists a payment function  $p: S \times [k] \to \mathbb{N}$  and a profile  $\pi$  such that  $\pi$  is a solution for the DNE problem over the payment game with the payment function *p*.

```
Yoav Feinstein, Orna Kupferman,
 Noam Shenwald
```

yoav.feinstein@mail.huji.ac.il, orna@cs.huji.ac.il, noam.shenwald@mail.huji.ac.il



## STREAM-BASED MONITORING OF ALGORITHMIC FAIRNESS

Jan Baumeister, Bernd Finkbeiner, Frederik Scheerer, Julian Siber, Tobias Wagenpfeil

### **UNFAIR AI SYSTEMS**



Al is used for critical decisions. COMPAS is an Al tool to predict recidivism. ProPublica revealed that **COMPAS** scores are biased against Afro-Americans.

### MONITORING FAIRNESS PROPERTIES



We propose to use the **RTLola monitoring framework** for **detecting unfairness** during the deployment of black-box AI systems through estimating conditional probabilities.

### SPECIFYING EQUALIZED ODDS WITH RTLOLA





## ERTIFYING PARETO OPTIMALITY IN MULTI-OBJECTIVE MAXSA

Christoph Jabs<sup>1</sup>, Jeremias Berg<sup>1</sup>, Bart Bogaerts<sup>2</sup>, Matti Järvisalo<sup>1</sup> <sup>1</sup>University of Helsinki, <sup>2</sup>KU Leuven & Vrije Universiteit Brussel

@ TACAS'25 Presentation Wednesday 11:30

#### **MULTI-OBJECTIVE OPTIMIZATION**

- Many real-world problems have multiple conflicting objectives.
- Aim: Pareto-optimal solutions
- Linear combination of objectives not sufficient

### - PROOF LOGGING -



- Write log of reasoning steps while solving
- Simple or formally verified checker
- Checking proof verifies correctness of result

[BOGAERTS ET AL. JAIR'23]

#### VERIPB

- Pseudo-boolean (cutting planes) proof system
- Derive constraints by linear combination
- Redundant constraints (RAT generalization [Järvisalo et al. IJCAR'12])
- Native single-objective support
- Preorder (originally for symmetry breaking), witness for redundant constraints must be smaller in preorder

### - MO-MAXSAT ALGORITHMS -

#### P-MINIMAL [SOH ET AL. CP'17]

- 1. Find a solution SAT solver
- Introduce PD cut
- Find a dominating solution If yes, goto 2
  - If no, previous is optimal, goto 1

#### Proof needs to cover

- SAT solver reasoning
- CNF Objective encodings
- PD cuts (see above)

### - EXPERIMENTAL EVALUATION -

### RESULTS

### PER-INSTANCE PROOF LOGGING OVERHEAD



#### MAXIMUM SATISFIABILITY

- Constraints: propositional formula, objective(s): linear function over variables
- Efficient for real-world optimization problems

### CONTRIBUTIONS

- VeriPB proofs for multi-objective MaxSAT
- No modifications to VeriPB proof system Open-source implementation for three
- algorithms
- Showing low overhead of proof logging

#### **MO-PROOFS**

Certificate that at least one representative solution for each non-dominated point was discovered.

#### PROOF SETUP

Unmodified VeriPB proof system

- 1. Encode Pareto dominance as preorder in proof
- Now Pareto-optimal solutions can only be explicitly excluded in the proof
- It the proof terminates as unsatisfiable, all Pareto-optimal solutions must appear in the proof

Syntactic restrictions

- First step in proof must load the Pareto order
- Order must never be changed

### PARETO DOMINANCE CUTS

- Blocking all solutions that are worse than the current one
- Similar to solution-improving constraint



Certifying a Pareto dominance cut for solution  $\alpha$ 

- Map each weakly dominated solution to  $\alpha$ ; Redundant with  $\alpha$  as witness
- 2. Exclude  $\alpha$  itself
- 3. Derive cut from steps 1 and 2

#### LOWER-BOUNDING

[CORTES ET AL. TACAS'23]

- Execute P-Minimal within upper-bounds on objectives
- Once done, loosen upper-bounds
- Upper-bounds on objectives can be ignored in proof  $\rightarrow$  same as *P*-Minimal

AVERAGE OVERHEAD AND CHECKING

w/logging w/o logging

1.233

1.220

1.247

 $\frac{\mathsf{checking}}{\mathsf{w}/\mathsf{logging}}$ 

47.52

21.81

29.70

#### BIOPTSAT [JABS ET AL. JAIR'24]

- Bi-objective
- Optimize first one objective, then the other
- When minimizing the second objective, PD ►
- cuts can be strengthened to unit clauses

#### In proof

- Derive lower-bound on first objective
- Certify PD cut
- Strengthen PD cut based on known lower-bound

### REFERENCES

Bart Bogaerts, Stephan Gocht, Ciaran McCreesh, and Jakob Nordström: Certified Dominance and Symmetry Breaking for Combinatorial Optimisation, JAIR, 2023.

João Cortes, Inês Lynce, and Vasco M. Manquinho: New Core-Guided and Hitting Set Algorithms for Multi-Objective Combinatorial Optimization, TACAS, 2023.

Christoph Jabs, Jeremias Berg, Andreas Niskanen, and Matti Järvisalo: From Single-Objective to Bi-Objective Maximum Satisfiability Solving, JAIR, 2024

Matti Järvisalo, Marijn Heule, and Armin Biere: Inprocessing Rules, IJ-CAR. 2012.

Takehide Soh, Mutsunori Banbara, Naoyuki Tamura, and Daniel Le Berre Solving Multiobjective Discrete Optimization Problems with Propositional Minimal Model Generation, CP, 2017.

> HELSINGIN YLIOPISTO HELSINGFORS UNIVERSITET UNIVERSITY OF HELSINKI





TIME

Algorithm

P-Minimal

BiOptSat

Lower-Bounding



## **Formally Verifying a Transformation from MLTL Formulas to Regular Expressions**



Zili Wang<sup>\*</sup>, Katherine Kosaian<sup>+</sup>, Kristin Yvonne Rozier<sup>\*</sup> \*Iowa State University, \*University of Iowa

[ziliw1, kyrozier]@iastate.edu, katherine-kosaian@uiowa.edu

### **Overview**

The results of formal verification are only as trustworthy as their input specifications. As such, the WEST tool was created to facilitate writing specifications in Mission-time Linear Temporal Logic (MLTL) by visualizing MLTL formulas as regular expressions (regexes). To certify the correctness of the WEST tool, we formally verify the correctness of the WEST algorithm in the interactive theorem prover Isabelle/HOL. We then generate a code export from our verified development in Haskell and use this to experimentally validate the existing WEST tool.



### **Mission-time LTL (MLTL)**

MLTL is a finite version of Linear Temporal Logic with discrete integer time bounds on the temporal operators. Here are some examples of traces (assignments of variables through time) that satisfy various MLTL formulas:

| <b>Operator</b>       | Syntax                       | 0           | 1   | 2          | 3            | 4          | 5        | 6            | 7   |
|-----------------------|------------------------------|-------------|-----|------------|--------------|------------|----------|--------------|-----|
| Globally              | ${\sf G}_{_{[2,5]}}p$        | $\bigcirc$  | •   | • <b>p</b> | p            | • p        | • p      | •            | •   |
| in the <b>F</b> uture | $\mathbf{F}_{[0,4]}p$        | $\bigcirc$  | •   | •          | $ \bigcirc $ | • p        | •        | •            | •   |
| Until                 | $p\mathbf{U}_{_{[1,6]}}q$    | $\bigcirc$  | • p | • p        | p            | • p        | <b>q</b> | •            |     |
| Release               | $p  \mathbf{R}_{_{[2,7]}} q$ | $ \bigcirc$ | •   | • <b>q</b> | q            | • <b>q</b> | q        | • <b>q</b> • | p,q |

### MLTL → Regular Expressions

Regular expressions (regexes) describe which variables must be true or false at certain timesteps. We say that a regex matches a trace. For example:

0, 1 S, 1 0, S • Fix a variable ordering match {**p**<sub>0</sub>, **p**<sub>1</sub>} ¬p  $\mathbf{p}_0$ ٦P, • 1 = true, 0 = false, p, S = any valueThe WEST algorithm: ss,s1,ss,ss,ss,ss,ss Transforms MLTL  $\rightarrow$  Regex Given an MLTL formula, ss.1s.s1.ss.ss.ss.ss computes the regular expressions ss,1s,1s,s1,ss,ss,ss which captures all the satisfying ss,1s,1s,1s,s1,ss,ss traces of the formula. For example on  $(\mathbf{p}_{\theta}\mathbf{U}_{[1,6]}\mathbf{p}_{1})$ , WEST produces the following: ss.1s.1s.1s.1s.s1.ss ss,1s,1s,1s,1s,1s,1s,s1 Check out all of our work here -

west.temporallogic.org



### **Formalization Contributions**

- Available on the Archive of Formal Proofs under the entry "Mission-time Linear Temporal Logic to Regular Expressions"
- We formalize each operation of the WEST algorithm in and prove overall correctness of the algorithm.
- We use Isabelle/HOL's code generator to obtain an implementation of the WEST algorithm in Haskell



assumes  $\pi$  long enough: "length  $\pi \geq$  complen mltl  $\varphi$ " shows "match  $\pi$  (WEST reg  $\varphi$ )  $\longleftrightarrow$  semantics mltl  $\pi \varphi$ "



Test suite: 1662 formulas that targets every combination of operator nesting Verified check: holds on 1658 formulas, times out on 4

Example - (¬p<sub>3</sub> U[0,2] ¬p<sub>2</sub>) R[0,2] (p<sub>0</sub> U[0,2] ¬p<sub>1</sub>)

Unverified check: Succeeded on all 1662



31st International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), May 5-8 2025, Hamilton, Canada

Acknowledgments: This work is supported by the NSF-GRFP 2024364991, NSF CAREER Award CNS-1552934, and NSF CCRI-2016592. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors(s) and do not necessarily reflect the views of the National Science Foundation.



### **INCREMENTAL SAT-BASED ENUMERATION OF** SOLUTIONS TO THE YANG-BAXTER EQUATION

Daimy Van Caudenberg<sup>1</sup>, Bart Bogaerts<sup>1,2</sup>,

### Leandro Vendramin<sup>3</sup>

<sup>1</sup>Declarative Languages and Artificial Intelligence section, KULeuven, Leuven, Belgium

<sup>2</sup>Department of Computer Science, Vrije Universiteit Brussel, Brussels, Belgium

<sup>3</sup>Department of Mathematics and Data Science, Vrije Universiteit Brussel, Brussels, Belgium



**KU LEUVEN** 

### THE YANG-BAXTER EQUATION

- Yang-Baxter equation (YBE) originally introduced in context of statistical [Yan67] and quantum mechanics [Bax72].
- Applications in knot theory, quantum computing...
- We focus on a subset of solutions.

decide if a formula is satisfiable.

- Finite, involutive, non-degenerate, combinatorial solutions.
- These solutions have relations to group and ring theory. • They can be studied using an equivalent mathematical structure: non-degenerate cycle sets.

SAT

• Boolean Satisfiability (SAT) problem decides if there exists a

• We encode the mathematical problem as a propositional for-

• hence, a satisfying assignment corresponds to a solution.

satisfying assignment for a given propositional formula.

• enumerate satisfying assignments (if they exist)...

### **ENUMERATION**

- Enumerating all (non-isomorphic) solutions to the YBE is an open problem!
- A database of solutions could inspire: • experimentation,
  - examples of algebraic structures to study,
- and counterexamples to previous conjectures...
- [AMV22] enumerated solutions up to size 10 using a constraint programming approach with static symmetry breaking.

### SAT MODULO SYMMETRIES

- SAT Modulo Symmetries (SMS) was originally introduced in [KS21,KS24]
- It forces a SAT solver to generate only non-isomorphic solutions during the search:
- Obtain a partial interpretation from the SAT solver.
- Perform a Minimality Check:
  - verify if the assignment can be extended to a complete assignment that is lexicographically minimal.
  - if this fails, the solver aborts the current branch of the search tree by learning a new clause.

**1. INTRO** 

<del>.</del>

### PARTIAL SOLUTIONS

- SMS introduces breaking constraints during the search, so, only a partial solution is known...
- This is an incomplete solution that can still be extended to a well-defined solution.
- A partial solution P is lexicographically minimal iff one of its extensions  $C \in \mathscr{X}(P)$  is lexicographically minimal,
- where  $\mathscr{X}(P)$  is the set of all complete solutions that P can be extended to

### **MINIMALITY CHECK FOR YBE**

- $\bullet$  Does there exist a permutation  $\pi$  s.t. the image of the current (partial) solution P under  $\pi$  is strictly smaller than P?
- If so, use  $\pi$  to exclude P and its extensions!
- Find breaking permutations by either:
- performing a backtracking search over the space of possible permutations,
- or by encoding as a SAT problem what it means for a permutation to be useful.



Enumeration Time per Size

10

AMV22 - Isomorphism Check Time Backtracking Approach - Minimality Check Time SAT-Based Approach - Minimality Check Time

### **FUTURE WORK** • Certifying the results.

- We obtain the same results as [AMV22], but that only means that we are either both correct or both wrong...
- Enumerating related structures
  - racks, used to enumerate skew cycle sets.
  - skew Cycle Sets, correspond to non-degenerate set-theoretic solutions.
  - biguandles, applications in knot theory.

[AMV22] Akgün, Ö., Mereb and M., Vendramin, L. (2022). Enumeration of Set-Theoretic Solutions to the Yang-Baxter Equation. In Mathematics of Computation 91(335). [YAM67] Yang, C. N. (1967). Some Exart Results for the Many-Body Problem in one Dimension with Repulsive Delta-Func-tion Interaction. In Phys. Rev. Lett. 19 (23). [BAX72] Baxter, R. J. (1972). Parition function of the Eight-Ventex lattice model. In Annals of Physics 70(1). [KS21] Kirchweger, M. and Szeider, S. (2021). SAT Modulo Symmetries for Graph Generation. In CP LIPIcs, vol. 210. [KS24] Kirchweger, M. and Szeider, S. (2024). SAT Modulo Symmetries for Graph Generation and Enumeration. In ACM Trans. Comput. Log., 25(3).

### Contact: Daimy.Vancaudenberg@kuleuven.be

### RESULTS

6 Problem Size

AMV22 - Total Time

Backtracking Approach - Total Time

Incremental Approach - Total Time

4

 $10^{6}$ 

 $10^{4}$ 

 $10^{2}$ 

 $10^{(}$ 

10

10

Time (s.)



• SAT solvers can:

mula.



Application-oriented Formal Verification Research Group Institute of Information Security and Dependability (KASTEL) Karlsruhe Institute of Technology

### **Revisiting Differential Verification: Equivalence Verification with Confidence**

Samuel Teuber | Philipp Kern | Marvin Janzen | Bernhard Beckert





For ReLU NNs  $f_1, f_2$ , input set  $X \subset \mathbb{R}^l$ and a *p*-norm  $\|\cdot\|_{p}$ :

**Equivalence Properties** 

 $\varepsilon$  equivalence [2, 3, 4]  $\forall x \in X. \|f_1(x) - f_2(x)\|_p \leq \varepsilon$ 

Top-1 equivalence [2]  $\forall x \in X. \operatorname{argmax}_{i} f_{1,i}(x) = \operatorname{argmax}_{i} f_{2,i}(x)$ 

### Classification Equivalence



Single NN verification via Zonotopes Propagate  $\mathcal{Z}^i \supseteq f^{(i)}(x)$ • Abstract Transformer for  $W \cdot \mathcal{Z} + \mathbf{b}$ 

Abstract Transformer for ReLU(Z)

**Differential Verification via Zonotopes** Also propagate  $\mathcal{Z}^i_{\Lambda} \supseteq f_1^{(i)}(x) - f_2^{(i)}(x)$ 

 $\mathcal{Z} = ((\mathbf{g}_1 \ \mathbf{g}_2 \ \mathbf{g}_3), \mathbf{c})$ 

### A Zonotope Transformer for ReLU Differences

 $\text{ReLU}(\mathcal{Z}_1) - \text{ReLU}(\mathcal{Z}_2) =$  $= \operatorname{ReLU}(\mathcal{Z}_1) - \operatorname{ReLU}(\mathcal{Z}_1 - \mathcal{Z}_{\Delta})$ 

Requires distinction over nine cases

Check max  $\leq$  0 for all  $j \neq k$ :

### Differential Verification for Classification



- Use differential bounds for LP formulation
- Verify classification equivalence
- Threshold T > 0: Confidence of  $f_1$
- $\max_{x} (\mathcal{Z}_{2}(x))_{j} (\mathcal{Z}_{2}(x))_{k}$ s.t.  $(\mathcal{Z}_{1}(x))_{l} + T \leq (\mathcal{Z}_{1}(x))_{k} \quad (l \neq k)$  $Z_{\Lambda} = Z_1 - Z_2$

### Linear Softmax Approximation





$$\left\{ \mathbf{z} \in \mathbb{R}^n \ \left| \ igwedge_{\substack{j=1 \ j 
eq i}}^n \mathbf{z}_i - \mathbf{z}_j \geq \ln\left(rac{\delta}{1-\delta}
ight) 
ight\}$$

$$\left\{ \mathbf{z} \in \mathbb{R}^n \mid \operatorname{softmax}{(\mathbf{z})_i \geq \delta} \right\}$$



Lack of asymmetry in Top-1 property

NP hard

Building on confidence-based verification [1]:

### $\delta$ -Top-1 equivalence

 $f_1, f_2$  have the same classification for every input where  $\operatorname{softmax}(f_1(x)) \geq \delta$ 

### **Evaluation**

Across 5 benchmarks, VeryDiff (our tool) outperforms other (non-)differential NN verification tools:

| Bon      | chmark      | Variant                 | Ι.  | Fouiv     | Counterey |           | Speedup |          |
|----------|-------------|-------------------------|-----|-----------|-----------|-----------|---------|----------|
| Ben      | ciinitii k  | v ui mine               |     | Equit.    |           | Junicical | Median  | Max      |
|          |             | VeryDiff (ours)         | 150 | (+24.0%)  | 153       | (+2.0%)   | _       | _        |
|          |             | NNEquiv                 | 37  |           | 142       |           | 37.3    | 8091.2   |
|          | ACAS        | MILPEquiv               | 16  |           | 3         |           | 7224.8  | 36297.0  |
| P .      |             | Marabou                 | 110 |           | 109       |           | 141.3   | 10070.5  |
| di p     |             | $\alpha, \beta$ -CROWN  | 121 |           | 150       |           | 15.4    | 1954.1   |
| εta      |             | VeryDiff (ours)         | 352 | (+101.1%) | 62        | (-43.6%)  | _       | -        |
| 01       | MNIST       | NNEquiv                 | 0   |           | 103       |           | 12.6    | 166.2    |
|          | (VeriPrune) | Marabou                 | 10  |           | 24        |           | 183.9   | 1390.8   |
|          |             | $\alpha, \beta$ -CROWN* | 175 |           | 110       |           | (516.9) | (4220.8) |
| H        | ACAS        | VeryDiff (ours)         | 169 | (+39.7%)  | 161       | (+103.8%) | _       | _        |
| Q 🕹      | ACAS        | NeuroDiff               | 121 |           | 79        |           | 43.1    | 16134.7  |
| s e      | MNIST       | VeryDiff (ours)         | 457 | (+11.5%)  | 242       | (+404.1%) | -       | -        |
| ž        | (VeriPrune) | NeuroDiff               | 410 |           | 48        |           | 4.5     | 1086.8   |
| 6 Tree 1 | LIIC        | VeryDiff (ours)         | 77  | (327.8%)  |           |           | -       | _        |
| 0-10p-1  | LIIC        | $\alpha, \beta$ -CROWN  | 18  |           | -         |           | 324.5   | 11274.3  |

### References

- [1] Anagha Athavale et al. "Verifying Global Two-Safety Properties in Neural Networks with Confidence". In: CAV'24. DOI: 10.1007/978-3-031-65630-9\_17.
- [2] Marko Kleine Büning et al. "Verifying Equivalence Properties of Neural Networks with ReLU Activation Functions" In: CP'20. DOI: 10.1007/978-3-030-58475-7\_50
- [3] Brandon Paulsen et al. "ReluDiff: differential verification of deep neural networks". In: ICSE'20. DOI: 10.1145/337 3380337
- [4] Samuel Teuber et al. "Geometric Path Enumeration for Equivalence Verification of Neural Networks". In: ICTAI'21. DOI: 10.1109/ICTAI52525.2021.00035

# Weakly acyclic diagrams: A data structure for infinite-

Michael Blondin<sup>1</sup>, Michaël Cadilhac<sup>2</sup>, **Xinyi Cui**<sup>3</sup>, Philipp Czerner<sup>3</sup>, Javier Esparza<sup>3</sup>, and **Jakob Schulz**<sup>3</sup>

### 1. Introduction

- Ordered Binary Decision Diagrams (OBDDs) [1, 2]
- $\diamond$  data structure representing fixedlength languages
- $\diamond$  efficient top-down dynamic programming for operations
- $\diamond$  used for finite-state symbolic model-checking

we generalize OBDDs to a class of infinite languages

### Weakly Acyclic Diagrams

- $\diamond$  data structure representing weakly acyclic languages
- $\diamond$  maintains algorithmic advantages of OBDDs
- $\diamond$  can be used for model checking in infinite-state systems

### 2. Weakly Acyclic Languages

- $\diamond$  Weakly acyclic languages: all cycles in DFA are self-loops
- $\diamond\,$  closed under union, intersection and complementation



### 3. Weakly Acyclic Diagrams

- data structure representing weakly acyclic languages using a table of nodes
- ◊ Node: state identifier, successors, accepting flag



### References

- S. Akers. 1978.
- [2] S. Chaki and A. Gurfinkel. 2018.
- [3] A. Finkel and Ph. Schnoebelen. 2001.
- [4] T. Geffroy, J. Leroux, and G. Sutre. 2017.
- $[5]\,$  A. Heußner, Gall T. L., and G. Sutre. 2009.
- [6] P. Ganty, C. Meuter, G. Delzanno, G. Kalyon, J.F. Raskin, and L. Van Begin. 2007.



### 5. Pre Operation

- $\diamond~$  Transducers: automata over the alphabet  $\Sigma\times\Sigma$  for fixed-length regular relations  $R\subseteq\Sigma^*\times\Sigma^*$
- $\diamond$   $\operatorname{Pre}_R(L) := \{ u \in \Sigma^* : (u, v) \in R, v \in L \}$  for a relation R and a language L



- $\diamond$  algorithm computing  ${\rm Pre}_R(L)$  given transducer for R and weakly acyclic DFA for L under assumption: language  ${\rm Pre}_R(L)$  is weakly acyclic
- $\diamond\,$  main idea: apply powerset construction on pairing of transducer and DFA contracting cycles in the process

1. mark 
$$M$$
  $pre(M = \{(p_0, q_0)\} \rightsquigarrow q_M$ 

2. 
$$M_a = \bigcup_{b \in \Sigma} \{(p', q') : (p, q) \in M, p \xrightarrow{(a,b)} p', q \xrightarrow{(b)} q'$$

3.  $s_M[a] = \begin{cases} \text{SELF} & \text{if } M_a \text{ is marked} \\ pre(M_a) & \text{otherwise} \end{cases}$ 

5. unmark M

4.  $b_M \iff \exists (p,q) \in M : p,q$  both accepting states

- $\diamond \ M \ {\rm set} \ {\rm of} \ {\rm state} \ {\rm pairs}, \ {\rm representing} \\ {\rm one} \ {\rm state} \ {\rm in} \ {\rm DFA} \ {\rm for} \ {\rm Pre}_R(L)$
- $\diamond\,$  use of markings to detect and remove cycles
- ◊ polynomial-time improvement if input transducer and DFA satisfy specific condition such that determinism is guaranteed

### 6. Application: Backwards Reachability

- $\diamond$  system as a set of configurations with a transition relation like lossy channel systems, Petri nets or broadcast protocols
- ◊ Backwards Reachability algorithm: computes the set of all predecessors of a given upwardclosed set of configurations, used for safety verification in systems [3]
- ♦ use Weakly Acyclic Diagrams for Backwards Reachability
  - \* set of configurations  $\rightsquigarrow$  weakly acyclic language
  - \* transitions  $\rightsquigarrow$  transducers
  - \* predecessor computation using pre

Left: Petri net, Right: transducer encoding transition t



### 7. Experimental Results

- library WADL for weakly acyclic diagrams implementing backwards reachability for lossy channel systems, Petri nets and broadcast protocols
- ◇ compared WADL to established safety verification tools BML[4], McScM[5] and MIST[6]
- $\diamond\,$  results show that WADL is competitive, solving most of the instances

| Benchmark         | BML    | McScM  | WADL  |
|-------------------|--------|--------|-------|
| tcp_simplest_err  | 0.04   | 0.01   | 0.04  |
| BAwCC_enh         | 4.17   | 124.72 | 0.50  |
| tcp_simplest      | 0.03   | 0.34   | 0.06  |
| peterson_3        | to     | to     | 19.50 |
| ring2             | 0.36   | SO     | 0.04  |
| brp_like_modified | 0.39   | 0.39   | 0.29  |
| simple_server     | 0.13   | 0.25   | 0.06  |
| рорЗ              | 400.20 | 4.73   | 2.07  |

SV-COMP'25 and Test-Comp'25 Posters



### 14<sup>th</sup> Competition on Software Verification

**Dirk Beyer** 



### **Participants**

Table 1: Competition candidates and representing jury members; Hors Concours participants are not listed since they are not represented by a jury member. <sup>new</sup> for first-time participants, <sup>meta</sup> for metaverifiers

| Participant            | Jury member          | Affiliation              |
|------------------------|----------------------|--------------------------|
| 2LS                    | V. Malík             | BUT, Czechia             |
| AISE                   | Z. Chen              | NUDT, China              |
| AProVE                 | N. Lommen            | RWTH Aachen, Germany     |
| BRICK                  | L. Bu                | Nanjing U., China        |
| Bubaak                 | M. Chalupa           | ISTA, Austria            |
| BUBAAK-SPLIT           | M. Chalupa           | ISTA, Austria            |
| COOPERACE meta ne      | WV. Vojdani          | U. Tartu, Estonia        |
| CPACHECKER             | M. Lingsch-Rosenfeld | LMU Munich, Germany      |
| CPV                    | PC. Chien            | LMU Munich, Germany      |
| DARTAGNAN              | H. Ponce de León     | Huawei Dresden, Germany  |
| DEAGLE                 | F. He                | Tsinghua U., China       |
| EmergenTheta           | L. Bajczi            | BME Budapest, Hungary    |
| ESBMC-INCR             | T. Wu                | U. Manchester, UK        |
| ESBMC-KIND             | T. Wu                | U. Manchester, UK        |
| GDART                  | F. Howar             | TU Dortmund, Germany     |
| Goblint                | S. Saan              | U. Tartu, Estonia        |
| HORNIX New             | M. Blicha            | U. Lugano, Switzerland   |
| JAVA-RANGER            | S. Hussein           | Ain Shams U., Egypt      |
| JBMC                   | P. Schrammel         | Diffblue, UK             |
| Korn                   | G. Ernst             | LMU Munich, Germany      |
| MLB                    | L. Bu                | Nanjing U., China        |
| Mopsa                  | R. Monat             | Inria & U. Lille, France |
| NACPA meta new         | H. Wachowitz         | LMU Munich, Germany      |
| Proton                 | R. Metta             | TCS, India               |
| RACERF <sup>new</sup>  | T. Dacík             | BUT, Czechia             |
| SVF-SVC <sup>new</sup> | M. Richards          | U. New South Wales, AU   |
| SV-SANITIZERS          | S. Saan              | U. Tartu, Estonia        |
| SWAT                   | N. Loose             | U. Luebeck, Germany      |
| Symbiotic              | M. Jonáš             | Masaryk U., Czechia      |
| Theta                  | L. Bajczi            | BME Budapest, Hungary    |
| THORN New              | L. Bajczi            | BME Budapest, Hungary    |
| UAUTOMIZER             | M. Heizmann          | U. Freiburg, Germany     |
| UGEMCUTTER             | D. Klumpp            | U. Freiburg, Germany     |
| UKojak                 | M. Bentele           | U. Freiburg, Germany     |
| UTAIPAN                | D. Dietsch           | U. Freiburg, Germany     |
|                        |                      |                          |

### **Cummulative Score**



Figure 1: Quantile plot for the category C-Overall.

### Ranking

Table 5: Overview of the top-three verifiers for each category; measurements for CPU time rounded to two significant digits.

| Rank     | Verifier                                      | Score | CPU<br>Time<br>(in h) | Solved<br>Tasks | Unconf.<br>Tasks | False<br>Alarms | Wrong<br>Proofs |
|----------|---|-------|-----------------------|-----------------|------------------|-----------------|-----------------|
| Reach    | Safety (11268 tasks, max. score 17860)        |       |                       |                 |                  |                 |                 |
| 1        | CPAchecker                                    | 10368 | 150                   | 6653            | 230              | 2               |                 |
| 2        | ESBMC-KIND                                    | 8717  | 69                    | 6830            | 599              |                 | 14              |
| 3        | CPV   | 7755  | 160                   | 6235            | 438              | 27              |                 |
| MemS     | afety (4042 tasks, max. score 6409)           |       |                       |                 |                  |                 |                 |
| 1        | CPAchecker                                    | 4892  | 18                    | 3818            | 1                |                 |                 |
| 2        | Symmotic                                      | 4479  | 2.3                   | 3671            | 0                |                 | 1               |
| 3        | UAUTOMIZER                                    | 3909  | 37                    | 2280            | 2                |                 | 1               |
| Concu    | rrencySafety (3175 tasks, max. score 5733)    |       |                       |                 |                  |                 |                 |
| 1        | Deagle  | 4604  | 3.1                   | 2500            | 38               | 1               | 4               |
| 2        | DARTAGNAN                                     | 3385  | 17                    | 2012            | 30               | 3               | 3               |
| 3        | UGEMCUTTER                                    | 3144  | 50                    | 1805            | 48               |                 |                 |
| NoOve    | rflows (8211 tasks, max. score 13297)         |       |                       |                 |                  |                 |                 |
| 1        | UAutomizer                                    | 11074 | 68                    | 6724            | 13               |                 |                 |
| 2        | UTAIPAN                                       | 10736 | 74                    | 6622            | 11               | 1               | 2               |
| 3        | UKOJAK  | 8878  | 54                    | 5910            | 2                |                 |                 |
| Termis   | nation (2328 tasks, max. score 4079)          |       |                       |                 |                  |                 |                 |
| 1        | Proton  | 3685  | 24                    | 1942            | 159              | 1               |                 |
| 2        | UAUTOMIZER                                    | 3334  | 18                    | 1667            | 4                |                 |                 |
| 3        | APROVE  | 2219  | 32                    | 1006            | 43               |                 |                 |
| Softwa   | reSystems (4329 tasks, max. score 7131)       |       |                       |                 |                  |                 |                 |
| 1        | CPAchecker                                    | 2178  | 30                    | 2022            | 55               |                 |                 |
| 2        | Mopsa   | 2086  | 20                    | 2212            | 0                |                 |                 |
| 3        | Symmotic                                      | 1822  | 6.1                   | 1487            | 231              |                 | 1               |
| Falsifie | cationOverall (30758 tasks, max. score 10675) |       |                       |                 |                  |                 |                 |
| 1        | CPAchecker                                    | 6999  | 100                   | 7100            | 67               | 2               |                 |
| 2        | Symmotic                                      | 6459  | 28                    | 6379            | 37               |                 |                 |
| 3        | BUBAAK  | 5565  | 18                    | 5739            | 236              | 9               |                 |
| Overal   | l (33353 tasks, max. score 55561)             |       |                       |                 |                  |                 |                 |
| 1        | UAutomizer                                    | 29710 | 270                   | 16677           | 196              |                 | 8               |
| 2        | CPACHECKER                                    | 26786 | 240                   | 20506           | 372              | 6               | 1               |
| 3        | Symmotic                                      | 20691 | 63                    | 14324           | 628              |                 | 3               |
| JavaO    | verall (673 tasks, max. score 926)            |       |                       |                 |                  |                 |                 |
| 1        | Java-Ranger                                   | 676   | 5.7                   | 491             | 13               |                 | 1               |
| 2        | JBMC  | 628   | 0.32                  | 430             | 91               |                 |                 |
| 3        | GDART   | 627   | 2.1                   | 460             | 15               |                 |                 |

## Features

Table 2: Algorithms and techniques that the participating verification systems used; "new for first-time participants,  $^{\varnothing}$  for hors-concours participation, and meta for meta-verifiers



### **Score Schema**

Table 6: Scoring schema for SV-COMP 2025 (unchanged from 2021)

| Reported result | Points  | Description   |
|-----------------|---------|---|
| UNKNOWN         | 0       | Failure to compute verification result  |
| FALSE correct   | $^{+1}$ | Violation of property in program was correctly found<br>and a validator confirmed the result based on a witness |
| False incorrect | -16     | Violation reported but property holds (false alarm)   |
| True correct    | +2      | Program correctly reported to satisfy property<br>and a validator confirmed the result based on a witness       |
| True incorrect  | -32     | Incorrect program reported as correct (wrong proof)   |

### Reference

#### Report

D. Beyer. State of the art in software verification and witness validation: SV-COMP 2025. In  $Proc.\ TACAS,$  LNCS . Springer, 2025

#### Acknowledgment

We thank the verification community for contributing their tools to the evaluation.

### **Frameworks**

Table 3: Solver libraries and frameworks that are used as components in the participating verification systems; <sup>new</sup> for first-time participants,  $^{\varnothing}$  for hors-concours participation and <sup>meta</sup> for meta-verifiers



### **Results**

Table 4: Quantitative overview over all regular results; empty cells are used for opt-outs, <sup>new</sup> for first-time participants, <sup> $\varnothing$ </sup> for hors-concours participation, and <sup>meta</sup> for meta-verifiers

| Participant                 | ReachSafety<br>11268 tasks<br>max. score 17860 | MemSafety<br>4042 tasks<br>max. score 6409 | ConcurrencySafety<br>3175 tasks<br>max. score 5733 | NoOverflows<br>8211 tasks<br>max. score 13297 | Termination<br>2228 tasks<br>max. score 4079 | SoftwareSystems<br>4329 tasks<br>max. score 7131 | FalsificationOverall<br>30758 tasks<br>max. score 10675 | Overall<br>33353 tasks<br>max. score 55561 | Ja va Overall<br>673 tasks<br>max. score 926 |
|-----------------------------|--|--|--|---|--|--|---|--|--|
| 2LS                         | 6053   | 686  | 0  | 6887  | 1703   | 1  | 1938  | 12658                                      |  |
| aise                        |  |  |  |   |  |  |   |  |  |
| AProVE                      |  |  |  |   | 2219   |  |   |  |  |
| BRICK                       | 0004   | 0.055                                      | 400  | 0580  | 4.404  | 4.008  |   | 4.88000                                    |  |
| Bubaak<br>Dubaak SaLis      | 6052   | 2240                                       | -190   | 6556  | 1491   | 1640   | 5569  | 16407                                      |  |
| CoOpeRacements rew          | 0000   | 3345                                       | =105   | 0330  | 1122   | 1045   | 3302  | 10431                                      |  |
| CPAchecker                  | 10368  | 4892                                       | 1770   | 8777  | 1301   | 2178   | 6999  | 26786                                      |  |
| CPV                         | 7755   |  |  |   |  |  |   |  |  |
| Dartagnan                   |  |  | 3385   |   |  |  |   |  |  |
| Deagle                      |  |  | 4604   |   |  |  |   |  |  |
| EmergenTheta                | 2106   |  |  | -361  | 620  |  |   |  |  |
| ESBMC-incr                  |  |  | 2155   |   |  |  |   |  |  |
| ESBMC-kind                  | 8717   | 3158                                       | 2155   | 8668  | 1115   | -1948  | 3741  | 18444                                      |  |
| Goblint                     | 2427   | 2198                                       | 2448   | 8486  | 969  | 545  |   | 17266                                      |  |
| Hornix                      |  |  |  |   |  |  |   |  |  |
| Monsa                       | 2807   | 2697                                       | 0  | 8491  | 0  | 2086   |   | 13521                                      |  |
| Nacpa <sup>meta new</sup>   | 10270  | 4887                                       | 1453   | 8843  | 1290   | 2127   | 6997  | 26131                                      |  |
| Proton                      |  |  |  |   | 3685   |  |   |  |  |
| RacerF                      |  |  |  |   |  |  |   |  |  |
| sv-sanitizers               |  | 861  |  | 1723  |  |  |   |  |  |
| SVF-SVC <sup>new</sup>      | -68717   | -10965                                     |  | -19469  |  | 0  |   |  |  |
| Symbiotic                   | 7097   | 4479                                       | 60   | 7704  | 1411   | 1822   | 6459  | 20691                                      |  |
| Theta                       | 3277   | 600  | 2275   | -170  | 712  |  |   |  |  |
| Thorn                       | -2519  | 2000                                       | 298  | -286  | 536  | 07.4   | 4700  | 20710                                      |  |
| UAutomizer                  | 3000   | 3909                                       | 2993   | 11074   | 3334   | 654  | 4762  | 29710                                      |  |
| UGenCutter                  | 4035   | 2030                                       | 3144   | 8878  | 0  | 300  | 3848  | 12872                                      |  |
| UTaipan                     | 6007   | 3711                                       | 2593   | 10736   | 0  | 288  | 4695  | 20244                                      |  |
| GDart                       |  |  |  |   |  |  |   |  | 627  |
| Java-Ranger                 |  |  |  |   |  |  |   |  | 676  |
| JBMC                        |  |  |  |   |  |  |   |  | 628  |
| MLB                         |  |  |  |   |  |  |   |  | 579  |
| SWAT                        | 4000   | 4008                                       | 04.0   | 8000  | 4400   | 0800   | 0504  | 0400                                       | 508  |
| CBMC <sup>2</sup>           | 1330   | 1885                                       | 819  | 7200  | 1199   | -2580  | -3581   | 9100                                       |  |
| CPA-bam-bnb <sup>2</sup>    |  | 3249                                       |  |   |  | -2370  |   |  |  |
| CPALockator <sup>20</sup>   |  | 3243                                       | -4967  |   |  | -4015  |   |  |  |
| Crux <sup>a</sup>           | 2133   |  |  | 608   |  |  |   |  |  |
| CSeq <sup>ø</sup>           |  |  | -12720   |   |  |  |   |  |  |
| DIVINE                      | 4629   | 502  | 351  | 0   | 0  | 72   | 352   | 3680                                       |  |
| EBF®                        |  |  | 360  |   |  |  |   |  |  |
| Frama-C-SV <sup>20</sup>    |  |  |  | 1573  |  |  |   |  |  |
| Gazer-Theta                 |  |  |  |   |  |  |   |  |  |
| GDart-LLVM"                 | 4041   |  |  |   |  | .670   | -820  | 4509                                       |  |
| Infor <sup>®</sup>          | -96489   |  | -8970  | -76213  |  | -32987   | -620  | 4008                                       |  |
| Lazy-CSeq <sup>37</sup>     | -00403   |  | -15153   | -10210  |  | 02001  |   |  |  |
| LF-checker <sup>ø</sup>     |  |  | 396  |   |  |  |   |  |  |
| Locksmith®                  |  |  |  |   |  |  |   |  |  |
| PeSCo-CPA <sup>S meta</sup> | 6269   |  |  |   |  | -1598  | 2435  | 16328                                      |  |
| PIChecker                   |  |  | 459  |   |  |  |   |  |  |
| Pinaka <sup>2</sup>         | 2448   | 1700                                       |  | 958   | 922  |  |   |  |  |
| PredatorHP <sup>20</sup>    | 11019  | 4733                                       |  |   |  |  |   |  |  |
| VeriAbs <sup>2</sup>        | 11012  |  |  |   |  |  |   |  |  |
| VeriOover <sup>®</sup>      | 11224  |  |  |   |  |  |   |  |  |
| WitnessMap                  |  |  |  |   |  |  |   |  |  |
| COASTAL                     |  |  |  |   |  |  |   |  | -3960  |
| JayHorn®                    |  |  |  |   |  |  |   |  | 248  |
| JDart <sup>o</sup>          |  |  |  |   |  |  |   |  | -1224  |
|                             |  |  |  |   |  |  |   |  |  |



## **CPV: A Circuit-Based Program Verifier**

### **Po-Chun Chien<sup>1</sup> and Nian-Ze Lee**<sup>2,1</sup>

<sup>1</sup>LMU Munich <sup>2</sup>National Taiwan University po-chun.chien@sosy.ifi.lmu.de nzlee@ntu.edu.tw



- **3rd** place in *ReachSafety-Overall*, including
- $\stackrel{\scriptstyle imes}{\scriptstyle o}$  **1st** in *ReachSafety-BitVectors* and *-ECA*
- 2nd in ReachSafety-Hardware and -ProductLines
- orginal States and Sta



### Summary

- Sequential circuits can serve as an intermediate representation for software verification.
- Offer different encoding options.
- Leverage powerful word- and bit-level hardware model checkers as backend.
- Perform competitively against established verifiers in SV-COMP.
- Ongoing work:
  - Support more verification properties (e.g., no-overflow and termination)
  - Export correctness witnesses
  - Apply circuit optimization to improve the performance of verification



### References

- [1] Beyer, D., Strejček, J.: Improvements in software ver-ification and witness validation: SV-COMP 2025. In: Proc. TACAS (3). LNCS 15698 (2025)
- [2] Biere, A., Froleyks, N., Preiner, M.: Hardware model checking competition 2024. In: Proc. FMCAD. pp. 7–7 (2024)
- [3] Biere, A., Heljanko, K., Wieringa, S.: AIGER 1.9 and beyond. Tech. Rep. 11/2, Institute for Formal Models and Verification, Johannes Kepler University (2011)
- Biere, A., Cimatti, A., Clarke, E.M., Strichman, O., Zhu, Y.: Bounded model checking. Advances in Computers
- **58**, 117–148 (2003) [5] Brayton, R., Mishchenko, A.: ABC: An academic industrial-strength verification tool. In: Proc. CAV. pp. 24–40. LNCS 6174 (2010)
- Chien, P.C., Lee, N.Z.: CPV: A circuit-based pro-gram verifier (competition contribution). In: Proc. TACAS (3). pp. 365–370. LNCS 14572 (2024)
- Cimatti, A., Griggio, A., Tonetta, S.: The VMT-LIB language and tools. In: Proc. SMT. CEUR Workshop Proceedings, vol. 3185, pp. 80–89 (2022) The VMT-LIB

- [8] Eén, N., Mishchenko, A., Brayton, R.K.: Efficient implementation of property directed reachability. In: Proc. FMCAD. pp. 125–134 (2011)

- [11] Griggio, A., Jonáš, M.: KRATOS2: An SMT-based model checker for imperative programs. In: Proc. CAV. pp. 423-436 (2023)
- [12] McMillan, K.L.: Interpolation and SAT-based model [13]
- Archman, R.E. Interpoteton and SAT-based model checking. In: Proc. CAV. pp. 1–13. LNCS 2725 (2003) Niemetz, A., Preiner, M., Wolf, C., Biere, A.: BTOR2, BTORMC, and BOOLECTOR 3.0. In: Proc. CAV. pp. 587– 595. LNCS 10981 (2018)
- Sheeran, M., Singh, S., Stâlmarck, G.: Checking safety properties using induction and a SAT-solver. In: Proc. FMCAD, pp. 127–144. LNCS 1954 (2000) [14]





### SV-COMP Benchmark: Verifying Intel TDX Module



### Dirk Beyer<sup>1</sup>, Po-Chun Chien<sup>1</sup>, Nian-Ze Lee<sup>2,1</sup>, & Thomas Lemberger<sup>1</sup>

<sup>1</sup>LMU Munich <sup>2</sup>National Taiwan University

### Intel Trust Domain Extensions



Source: Fig. 2-1 in Intel TDX Module v1.5 Base Arch. Spec. [2]

### Defining Verification Tasks

```
name: tdh_mng_create__requirement__expected
     target:
       filename: formal/harness/tdh_mng_create_harness.c
       method: tdh_mng_create__valid_entry
     before target:
        filename: formal/src/initialization.c
         method: init_tdx_general
         filename: formal/src/initialization.c
         method: init_vmm_dispatcher
10
       - filename: formal/harness/tdh_mng_create_harness.c
         method: tdh_mng_create__common_precond
     after_target:
12
        - filename: formal/harness/tdh_mng_create_harness.c
13
14
         method: tdh_mng_create__common_postcond
15
     properties:
        property_file: unreach-call.prp
16
         expected_verdict: true
17
```

### Nondet Initialization of struct

Havoc memory: by assigning a nondeterministic value to each byte

Havoc object: by nondeterministically initializing each field of the type (if a field is a non-primitive type, recursively initialize it)

Verifier builtin: e.g., \_\_CPROVER\_havoc\_object in CBMC

```
_NONDET_struct_tdvps_t(tdvps_t* dest) {
  void
    _NONDET_custom_type(dest, sizeof(tdvps_t));
  }
3
  void _NONDET_custom_type(void* base, unsigned int size) {
4
    for (int i
               = 0; i
                       < size; i++
      *((char*)base + i) = _NONDET_uint8t();
7
 }
               Initialization by havocking memory
      _NONDET_struct_tdvps_s(struct tdvps_s *dest) {
void
   _NONDET_struct_tdvps_ve_info_s(&((*dest).ve_info));
```

```
NONDET_array_1D_unsigned_char(&((*dest).reserved_0), 128);
// ... snipped ...
}
```

```
void _NONDET_array_1D_unsigned_char(unsigned char (*dest)[],
vint dim0) {
```

```
s for (int i = 0; i < dim0; i++)
9 (*dest)[i] = _NONDET_uchar();
10 }</pre>
```

Initialization by havocking object

### References

- Intel TDX Module v1.5 ABI Specification, https://www.intel.com/content/www/us/en/ content-details/795475/intel-tdx-module-v1-5-abi-specification.html, accessed: 2024-05-01
- [2] Intel Trust Domain Extensions, https://www.intel.com/content/www/us/en/developer/ tools/trust-domain-extensions/documentation.html, accessed: 2024-05-01

### ABI Specifications

- TDs and VMM communicate through Application Binary Interfaces
- **Goal**: Verify TDX ABIs (implemented in C + assembly) adhere to the specification under all inputs

#### Table 5.145: TDH.MNG.CREATE Input Operands Definition

| Operand | Description  |                |   |
|---------|--|----------------|---|
| RAX     | SEAMCALL instruction leaf number and version, see 5.3.1                        |                |   |
|         | Bits   | Field          | Description                                     |
|         | 15:0   | Leaf Number    | Selects the SEAMCALL interface function         |
|         | 23:16  | Version Number | Selects the SEAMCALL interface function version |
|         |  |                | Must be 0                                       |
|         | 63:24  | Reserved       | Must be 0                                       |
| RCX     | The physical address of a page where TDR will be created (HKID bits must be 0) |                |   |
| RDX     | RDX Bits Name Description  |                | Description                                     |
|         | 15:0   | HKID           | The TD's ephemeral private HKID                 |
|         | 63:16  | Reserved       | Reserved: must be 0                             |

#### Table 5.146: TDH.MNG.CREATE Output Operands Definition

| Operand | Description                                  |  |
|---------|--|--|
| RAX     | SEAMCALL instruction return code - see 5.3.1 |  |
| Other   | Unmodified                                   |  |

Example: Specification of ABI TDH.MNG.CREATE [1]

### Verification Harnesses

- Initialize global data and assume preconditions
- Mock access to externally defined data and model inline assembly
- Check postconditions

### Contributions

- 290 tasks from 16 host-side ABIs (TDH) and 5 guest-side ABIs (TDG)
- *F*HARNESSFORGE (gitlab.com/sosy-lab/software/harnessforge):
  - Assembles single-file verification tasks from real-world C projects
  - Slices off code irrelevant to verification tasks
- Next steps:
  - Experiment with effect of different initialization strategies
  - Develop more tooling to support harness generation (e.g., harnessspecific linter)
  - Establish custom annotations for initializing complex types in SV-COMP community

#### More Information

Intel TDX Module Verification tasks







This work is supported by a research gift from Intel.

### AProVE (KoAT + LoAT)

### Automatic Termination Analysis of C Programs

Nils Lommen, Florian Frohn, and Jürgen Giesl

### Overview

- AProVE (KoAT + LoAT) [1] is a framework to analyze termination of C Programs
- Programs are transformed into Integer Transition Systems (ITSs)
- ITSs are analyzed by our tools KoAT [2] and LoAT [3]



### **LLVM Program**

- C program is compiled into LLVM code using Clang.
- LLVM fragment of the loop body:

| 1 | %10 = load %1           | # load ×             |
|---|-------------------------|----------------------|
| 2 | $\%11 = mul \ 3 \ \%10$ | # multiply x by 3    |
| 3 | store %11, %1           | # store x            |
| 4 | %12 = load %2           | # load y             |
| 5 | $\%13 = mul \ 2 \ \%13$ | # multiply y by 2    |
| 6 | store %13, %2           | # store y            |
| 7 | br %6                   | # jump to loop guard |
|   |                         |                      |

### Symbolic Execution Graph (SEG) & ITS

SEG represents all possible program runs, augmented with invariants:

- Its nodes are *abstract states* that represent sets of actual program states
- SEG handles the heap, pointer arithmetic, and recursive data structures
- $\bullet~\mbox{LLVM}$  code is transformed automatically into an SEG

ITSs are a simple language for integer programs:

- Turing-complete formalism with only integer variables over  $\ensuremath{\mathbb{Z}}$
- SEG is transformed into ITS

## $\rightarrow \underbrace{\ell_0}_{t_0} \underbrace{\ell_1}_{\eta(x)} \underbrace{t_1 : \varphi = (x < y)}_{\eta(x)} \\ \eta(x) = 3 \cdot x \\ \eta(y) = 2 \cdot y$

### KoAT (Termination & Upper Time Bounds)

- Automated complexity and termination analysis of ITSs
- Alternating modular inference of runtime and size bounds
- How often can a transition be executed?
  - Multiphase Linear Ranking Functions
     → Use SMT-solver Z3 to infer well-founded relation

TWN-Loops
 → Reduce termination problem to SMT problem
 Completeness for the class of so-called TWN-loops

- How *large* are the variables?
  - Compute bounds for each change of a variable  $\hookrightarrow$  Over-approximate the number of changes by runtime bounds
  - Use runtime bounds and closed forms of loops

## 

### LoAT (Non-Termination and more)

### Features

- Techniques
- non-termination ADCL DFS + acceleration
- *lower* time bounds **ABMC** BFS + acceleration
- safety / unsafety **TRL** BFS + recurrence analysis

### Non-term. via <u>A</u>cceleration <u>D</u>riven <u>C</u>lause <u>L</u>earning

- Depth-first exploration of state space
- Applies <u>acceleration</u> when a loop is encountered under-approximation of the loop's transitive closure
- Non-term. proofs as "by-product" of acceleration
- Exploits *redundancy* to cut off infinite branches



### References

- [1] Nils Lommen and Jürgen Giesl. AProVE (KoAT + LoAT) Website: https://koat.verify.rwth-aachen.de/svcomp25.
- [2] Nils Lommen, Éléanore Meyer, and Jürgen Giesl. KoAT Website: https://koat.verify.rwth-aachen.de/.
- [3] Florian Frohn and Jürgen Giesl. LoAT Website: https://loat-developers.github.io/LoAT/.





### Exemplary C Program

int main() {

Does the following program terminate?

extern int \_nondet(void);

int x = \_nondet();

int y = \_nondet();

while (x < y) { x = 3\*x;

y = 2\*y;

return 0;



## PAchecker A Tool for Configurable Program Analysis



Daniel Baier, Dirk Beyer, Marek Jankola, Matthias Kettl, Marian Lingsch-Rosenfeld, and Philipp Wendler

### CPACHECKER

CPACHECKER is a modern and versatile framework for building software-verification analyses from well-known concepts that match the user's requirements.





### VERIFICATION

For SV-COMP 2025, we used CPACHECKER 4.0 with strategy selection to choose a parallel portfolio of analyses suitable for a given verification task. More details can be found in our tutorial [1].

- Strategy selection chooses a *parallel* portfolio of analyses. The parallel composition is denoted by the || symbol
- Support all properties and categories of C programs
- 1st place in the categories FalsificationOverall, ReachSafety, MemorySafety, and SoftwareSystems
- 2nd place in the category Overall
- Only 7 wrong results out of  $33\,353$  tasks  $(0.02\,\%)$
- New and improved analyses for:
  - Reachability
  - Memory safety
  - Termination



Paper [1] available here

### WITNESS VALIDATION

For witness validation in SV-COMP 2025, we used CPACHECKER 4.0 with strategy selection to choose (based on the input) a mature analysis. More details can be found in our competition contribution [5].

- Strategy selection chooses a mature analysis that is suitable for a given validation task
- Support all witness types and formats, all properties for which witnesses exist, and all categories of C programs
- 1st place in the categories MemorySafety, Termination, and SoftwareSystems for violation witnesses 1.0
- 2nd place in all categories for correctness witnesses 2.0, most categories for correctness witnesses 1.0, and most categories for violation witnesses 2.0



Paper [5] available here

Contributors

### VERIFICATION STRATEGY FOR SV-COMP 2025



CPACHECKER is an open-source project, mainly developed by the Software and Computational Systems Lab at LMU Munich, and is used and extended by international associates from U Passau, U Oldenburg, U Paderborn, ISP RAS, TU Prague, TU Vienna, TU Darmstadt, and VERIMAG in Grenoble, along with several other universities and institutes.



### References

- Baier, D., Beyer, D., Chien, P.C., Jakobs, M.C., Jankola, M., Kettl, M., Lee, N.Z., Lemberger, T., Lingsch-Rosenfeld, M., Wachowitz, H., Wendler, P.: Software verification with CPACHECKER 3.0: Tutorial and user guide. In: Proc. FM. pp. 543-570. LNCS 14934, Springer (2024). https://doi.org/10.1007/978-3-031-71177-0\_30
   Beyer, D., Dangl, M., Wendler, P.: Boosting k-induction with continuously-refined invariants. In: Proc. CAV. pp. 622-640. LNCS 9206, Springer (2015). https://doi.org/10.1007/978-3-319-21690-4\_42
   Beyer, D., Dangl, M., Wendler, P.: A unifying view on SMT-based software verification. J. Autom. Reasoning 60(3), 299-335 (2018). https://doi.org/10.1007/s10817-017-9432-6
   Beyer, D., Lee, N.Z., Wendler, P.: Interpolation and SAT-based model checking revisited: Adoption to software verification. J. Autom. Reasoning 69 (2025). https://doi.org/10.1007/s10817-024-09702-9, preprint:

https://doi.org/10.48550/arXiv.2208.05046

- https://doi.org/10.48550/arXiv.2208.05046
  [5] Beyer, D., Lingsch-Rosenfeld, M.: CPACHECKER VALIDATOR 4.0 (competition contribution). In: Proc. TACAS (3). LNCS 15698, Springer (2025)
  [6] Beyer, D., Löwe, S.: Explicit-state software model checking based on CEGAR and interpolation. In: Proc. FASE. pp. 146–162. LNCS 7793, Springer (2013). https://doi.org/10.1007/978-3-642-37057-1\_11
  [7] Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement for symbolic model checking. J. ACM 50(5), 752–794 (2003). https://doi.org/10.1145/876638.876643
  [8] Friedherger, K.: CPA-BAW: Block-abstraction memoization with value analysis and model.
- Friedberger, K.: CPA-BAN: Block-abstraction memoization with value analysis and pred-icate analysis (competition contribution). In: Proc. TACAS. pp. 912–915. LNCS 9636, Springer (2016). https://doi.org/10.1007/978-3-662-49674-9\_58



## **O**theta

### Modular, Highly Configurable, Automatic Model Checker





Theta · Budapest, Hungary · Critical Systems Research Group · https://ftsrg.mit.bme.hu



### 7<sup>th</sup> Competition on Software Testing Dirk Beyer



LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

### **Features**

Table 2: Technologies and features that the test generators used



### **Results**

Table 3: Quantitative overview over all results

| Participant                                     | <b>Cover-Error</b><br>1215 tasks | <b>Cover-Branches</b><br>10011 tasks | <b>Overall</b><br>11226 tasks |
|---|----------------------------------|--------------------------------------|-------------------------------|
| $\mathbf{Cetfuzz}^{\varnothing}$                | 323                              | 2524                                 | 2906                          |
| CoVeriTest                                      | 552                              | 4959                                 | 5333                          |
| ESBMC-incr                                      | 679                              | 4380                                 | 5591                          |
| ESBMC-kind                                      | 680                              | 4323                                 | 5565                          |
| FDSE  | 729                              | 5468                                 | 6435                          |
| Fizzer  | 736                              | 5429                                 | 6446                          |
| FuSeBMC   | 994                              | 5656                                 | 7763                          |
| $\mathbf{FuSeBMC-AI}^{\varnothing}$             | 853                              | 4077                                 | 6228                          |
| $\mathbf{HybridTiger}^{\varnothing}$            | 438                              | 3866                                 | 4193                          |
| $\mathbf{KLEE}^{\varnothing}$                   | 804                              | 3065                                 | 5434                          |
| KLEEF   | 969                              | 5734                                 | 7692                          |
| $\mathbf{Owi}^{\varnothing}$                    | 281                              | 2462                                 | 2677                          |
| PRTest  | 211                              | 3191                                 | 2764                          |
| $\mathbf{Rizzer}^{\varnothing}$                 | 608                              |                                      |                               |
| Sikraken <sup>new</sup>                         |                                  | 2469                                 |                               |
| Symbiotic                                       | 743                              | 4207                                 | 5793                          |
| TracerX   | 390                              | 3327                                 | 3667                          |
| TracerX-WP                                      | 349                              | 3275                                 | 3447                          |
| UTestGen  | 439                              | 4393                                 | 4492                          |
| $\mathbf{WASP}\text{-}\mathbf{C}^{\varnothing}$ | 554                              | 2740                                 | 4094                          |

### References

### Reference

D. Beyer. Automatic testing of C programs: Test-Comp 2025. Springer, 2025

### Funding

This project was funded in part by the Deutsche Forschungsgemeinschaft (DFG) — 418257054 (Coop).

Table 1: Competition candidates with tool references and representing jury members; <sup>new</sup> indicates first-time participants

**Participants** 

| Tester                   | License | Jury member    | Affiliation                          |
|--------------------------|---------|----------------|--------------------------------------|
| Cetfuzzø                 | Apache  | _              | _                                    |
| CoVeriTest               | Apache  | MC. Jakobs     | LMU Munich, Germany                  |
| ESBMC-incr               | Apache  | C. Wei         | U. of Manchester, UK                 |
| ESBMC-KIND               | Apache  | C. Wei         | U. of Manchester, UK                 |
| FDSE                     | Apache  | Z. Chen        | National U. Defense Techn., China    |
| Fizzer                   | Zlib    | M. Trtík       | Masaryk U., Brno, Czechia            |
| FUSEBMC                  | MIT     | K. Alshmrany   | U. of Manchester, UK and             |
|                          |         |                | Inst. Public Admin., Saudi Arabia    |
| FUSEBMC-AI               | MIT     | _              | _                                    |
| HybridTiger <sup>2</sup> | Apache  | _              | _                                    |
| KLEEF                    | NCSA    | A. Misonizhnik | Independent Researcher, Neutral      |
| KLEE <sup>Ø</sup>        | NCSA    | _              | _                                    |
| Owı∞                     | AGPL    | _              | _                                    |
| PRTEST                   | Apache  | T. Lemberger   | LMU Munich, Germany                  |
| Rizzer                   | Zlib    | _              | _                                    |
| SIKRAKEN <sup>new</sup>  | LGPL    | C. Meudec      | South East Technological U., Ireland |
| Symbiotic                | MIT     | M. Jonáš       | Masaryk U., Brno, Czechia            |
| TRACERX                  | Apache  | J. Jaffar      | National U. of Singapore, Singapore  |
| TRACERX-WP               | Apache  | J. Jaffar      | National U. of Singapore, Singapore  |
| UTESTGEN                 | LGPL    | M. Barth       | LMU Munich, Germany                  |
| $WASP-C^{\varnothing}$   | Apache  | _              | _                                    |
| TestCov                  | Apache  | M. Kettl       | LMU Munich, Germany                  |

### **Final Score**

Figure 1: Quantile functions for category Overall.







Ranking

Table 4: Overview of the top-three test generators for each category (measurement values for CPU time in hours, rounded to two significant digits)

| Rank           | Tester      | Score | CPU             |  |  |  |
|----------------|-------------|-------|-----------------|--|--|--|
|                |             |       | $\mathbf{Time}$ |  |  |  |
| Cover-E        | Cover-Error |       |                 |  |  |  |
| 1              | FuSeBMC     | 994   | 75              |  |  |  |
| 2              | KLEEF       | 969   | 9.5             |  |  |  |
| 3              | Symbiotic   | 743   | 5.5             |  |  |  |
| Cover-Branches |             |       |                 |  |  |  |
| 1              | KLEEF       | 5734  | 1500            |  |  |  |
| 2              | FUSEBMC     | 5656  | 2500            |  |  |  |
| 3              | FDSE        | 5468  | 2200            |  |  |  |
| Overall        |             |       |                 |  |  |  |
| 1              | FuSeBMC     | 7763  | 2600            |  |  |  |
| 2              | KLEEF       | 7692  | 1500            |  |  |  |
| 3              | Fizzer      | 6446  | 2100            |  |  |  |



**Coverifier-Based Testing** 

LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Marie-Christine Jakobs (m.jakobs@lmu.de)



D. Beyer and S. Löwe. Explicit-state software model checking based on CEGAR and interpolation. In Proc. FASE, LNCS 7793, pages 146–162. Springer, 2013.
 D. Beyer, M. E. Keremoglu, and P. Wendler. Predicate abstraction with adjustable-block encoding. In Proc. FMCAD, pages 189–197. FMCAD, 2010.



### TracerX: Enhancing Dynamic Symbolic Execution with Weakest Precondition & **Deletion Interpolation**

Arpita Dutta<sup>1</sup>, Rasool Maghareh<sup>2</sup>, Joxan Jaffar<sup>3</sup>, Sangharatna Godboley<sup>4</sup>, and Xiao Liang Yu<sup>3</sup>

[arpita, joxan, xiaoly]@comp.nus.edu.sg<sup>1,3,5</sup>, sajjadrsm@gmail.com², sanghu@nitw.ac.in<sup>4</sup> Iniversity of Singapore, Singapore<sup>1,3,5</sup>, Lemurian Labs, Toronto, Canada², National Institute of Technology Warangal, India<sup>4</sup> School of Computing, National Univ



#### Highlights

- TracerX is Dynamic Symbolic Executor (DSE) which is built upon KLEE [1].
- The major advantage of DSE is its path-by-path exploration of the program execution space. However, this often leads to the path explosion problem.
- To address this issue, a method of abstraction learning has been used. The key step here is the computation of an interpolant to represent the learned abstraction [2].
- We use two different approaches of interpolant generation viz., 1. Deletion Interpolation [3] and 2. Weakest Precondition (WP) Interpolation [4]
- Deletion Interpolation (TracerX-Del) is our more stable and mature system and is briefly discussed in [3].
- Here, we present the Weakest Precondition (WP) Interpolation approach for TracerX, i.e., TracerX-WP

#### From KLEE (No Interpolation) to TracerX (With Interpolation)

- Forward Symbolic Execution to find feasible paths (Similar to KLEE).
- Intermediate execution states preserved (Unlike KLEE).
- Half interpolant aka (PATH Interpolants) are generated during backward tracking and Full interpolants aka (TREE Interpolants) are generated by merging the half interpolants.
- Full interpolants used for subsumption at similar program points.
- TracerX uses information from already traversed subtree to prune other subtrees.



Figure 2. Exploration of Symbolic Execution Tree in Non-pruning DSE vs. Pruning DSE

Symbolic Execution Tree with Interpolation



#### Interpolation: Weakest Precondition

- Ideal interpolant is the weakest precondition (WP) of the target. Unfortunately, WP is intractable to compute, which means it is difficult or impossible to find an exact solution.
- For example, in the above code snippet: assume  $(not(b1 \land \neg b2 \land \neg b3))$ .
- Hence, the WP before the first if-statement is:
- WP is:  $b1 \longrightarrow (\neg b2 \land b3 \land x \le 7) \lor (b2 \land x \le 4)$  $b1 \longrightarrow x < 3$
- Essentially, WP is exponentially disjunctive. This means that any one of the conditions can be satisfied for the target to be reached. (As shown here)
- One way to approximate WP is to use a conjunctive approximation, which involves expressing the WP as a conjunction of simpler conditions (Challenge is to obtain a conjunctive approximation).

### Approximation of Weakest Precondition

A Path is a sequence of assignments and assume instructions.

- 1. Interpolant of Assignment instruction:
  - wp( $inst, \omega$ ) = · · · inverse transition of inst over  $\omega$  Implemented at LLVM IR level: LD/ST, add, sub, cmp, cast, GEP, etc.
  - e.g.  $\omega : x \le 15$  and inst : x = z + 2, then  $wp(inst, \omega) : z \le 13$
- 2. Interpolant of Assume instruction (C is incoming Context):
  - $\{C\}$  assume(B)  $\{\omega\}$
  - WP Approximation: find  $\bar{C}$  to replace C
     ABDUCTION PROBLEM !!!

#### Following algorithm is the heart of TracerX:

1. We compute finest partition so that  $var(C_i) * var(C_j) s.t. i \neq j$ :  $\{C_1 * C_2 * C_3 * ... * C_n\}$  assume(B)  $\{\omega_1 * \omega_2 * \omega_3 * \dots * \omega_m\}$  (\* is as in separation logic).

#### 2. Bunch $C_i$ into three:

- Target independent: The C<sub>i</sub> which are separate from B and ω.
   Action: Replace C<sub>i</sub> with true, i.e. remove C<sub>i</sub>.
- Guard independent: Consider C<sub>gi</sub> ≡ C<sub>i</sub> s.t. C<sub>i</sub> \* B; and, ω<sub>gi</sub> ≡ ω<sub>j</sub> s.t. B \* ω<sub>j</sub>.
- ction: Replace  $C_{gi}$  by  $\omega_{gi}$
- Remainder of the C<sub>i</sub>: We do not capture exact WP for this group.
- Action: No change to  $C_i$ , i.e. keep  $C_i$ .

#### KLEE [1] v/s TracerX-Del [3] v/s TracerX-WP [4]

KLEE

TracerX-WP

TracerX-Del

TracerX (Del/WP)

80

20 25

Value of N

Consider this program; here, function f(N), returns the sum of the first N natural numbers.



- timeout for N=17. Both TracerX-Del and TracerX-WP scale up to
- N=50 [timeout=600 secs for each tool] Clear gap in the number of subsumed paths
- between TracerX-Del and TracerX-WP
- TracerX-WP generates better interpolants than TracerX-Del for subsumption.
- Reason: TracerX-Del generates interpolants as a subset of the incoming context whereas TracerX-WP generates interpolants from the weakest precondition of a path.

#### **Experimental Results**

aths

ned

Data set: All C-programs from RERS-2012 Challenge [6].

Total targets: 1159

All three systems KLEE [1], CBMC [5] and TracerX-WP [4] are run for 3600 seconds

#### Observations

- 1. TracerX-WP able to detect 348 targets, while KLEE and CBMC are detected 245 and 117 targets respectively.
- 2. All the targets reached by TracerX-WP are super set of targets covered by CBMC and KLEE.
- 3. TracerX-WP is 29.59x faster than KLEE and 66.37x faster than CBMC.



#### References

- [1] C. Cadar et al. Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs. In: OSDI, 2008
- [2] J. Jaffar et al. TRACER: A symbolic execution tool for verification. In: CAV, 2012.
- [3] J. Jaffar et al. TracerX: Dynamic symbolic execution with interpolation (competition contribution). In: FASE, 2020.
- [4] A. Dutta et al. TracerX: Pruning Dynamic Symbolic Execution with Deletion and Weakest Precondition Interpolation (competition contribution). In: FASE, 2024
- [5] D. Kroening D et al. CBMC-C Bounded Model Checker. In: TACAS 2014.
- [6] http://rers-challenge.org/

