muRelBench: MicroBenchmarking for Zonotope Domains

Kenny Ballou[0000−0002−6032−474X] and Elena Sherman[0000−0003−4522−9725]

Boise State University kennyballou@u.boisestate.edu,elenasherman@boisestate.edu

Abstract. We present `muRelBench`, a suite of synthetic benchmarks for weakly-relational abstract domains and their operations. For example, the benchmarks can support experimental evaluations of proposed algorithms such as domain closure.

Keywords: Weakly-Relational Abstract Domains · Benchmarks · Tests

1 Introduction

Zonotopes [9], relational numerical abstract domains, are widely used in program and system verification using static analysis and model-checking techniques and, recently, found their way into the verification of neural networks [10]. To reason about their computations, verifiers manipulate abstract domains through a predefined set of operations, e.g., Least-Upper Bound (LUB), closure, or forget operators [14]. Such manipulations of abstract states commonly dominate the computation time of a verifier. Thus, there has been extensive research on improving the efficiency of operations of Zonotopes such as closure [1, 3, 7, 14, 17].

While the new algorithms provide their complexity estimates, empirically evaluating their results remains crucial to comprehensively assessing their impact. Commonly, such evaluations are performed in the context of a verifier and its target, e.g., a data-flow analyzer using Zones [12, 13] on a set of programs. However, depending on program structure and semantics [2], one may or may not detect the effect of the new operation on Zones. Thus, the question becomes whether the set of programs is not representative or the implementation of the new algorithm is inefficient and requires additional tuning. Because of the complexity of Zonotope states, it is difficult to assess whether a verifier produces states with properties that a novel operation algorithm takes an advantage of.

This problem is known to other research communities such as software engineering and compiler optimization community, which they solv by establishing microbenchmarking frameworks [11]. Microbenchmarking isolates the effects of a specific technique such as a certain optimization on syntactically generated code with desired features. In this work, we introduce an extensible `muRelBench` microbenchmarking framework for Zonotopes that is built on top of the `JMH` [15,16] profiling tool for Java programs. `muRelBench` eliminates verifier and program dependencies and focuses on specific operations of parameterized Zonotope states.

For a given type of Zonotope domain, Z and its operation *ops*, `muRelBench` takes as an input set of predefined parameters for each characteristic of the corresponding Z typed abstract domain. Then the framework exhaustively generates abstract states corresponding to each element of the Cartesian product of those parameters and applies *ops* and correctness checks, if any, within the `JMH` context. Upon the completion of experiments, `muRelBench` writes the runtime results for each abstract domain to a Comma-Separated Values (CSV) file, that researchers can use for further evaluations.

In its current version, generation of abstract states is parameterized by the number of variables and variable connectivity for Octagon(Z) [14]. Thus, synthetically generated matrices that encode Octagon states vary in their size and density. `muRelBench` implements two closure operations (*ops*): Full Transitive Closure (using Floyd-Warshall all pairs shortest path [5]) and Chawdhary [3] incremental closure. However, as we describe in the next section, `muRelBench` can be easily extended to different Z and *ops* types.

Thus, this microbenchmarking framework has the following three key features: (1) dynamic generation of parameterized abstract states, (2) application of user defined operations on them, and (3) checks to user-defined properties, e.g., pre/post conditions on Zonotope states before and after executing operations. We believe that `muRelBench` will help rapid prototyping of abstract operations and evaluating the efficiency of existing implementations.

In the next section 2, we describe framework details and explain how it generates different abstract states. To demonstrate the usefulness of `muRelBench`, in Section 3, we present a case study on runtime data of the two closure operators on Octagon states. We conclude the paper with future work on `muRelBench`.

2 `muRelBench` Framework

Figure 1 provides overview of `muRelBench`'s components. In the dashed rounded rectangle are user-defined components of an abstract domain type Z, operations, e.g., *ops*1, and property checks of the state after *ops*1 modifies the abstract state. These bindings are defined at compile-time. *State generator* component takes generation parameters N and C, and Z type, and randomly (up to the seed) generates $N \times C$ abstract states.

Benchmarking component takes the generated states and applies *ops*1 state operation and checks the results with *check*1. The component also takes the runtime parameters for `JMH` that defines what type or runtime data to collect and how many times to repeat the experiments. Upon completion, the data is written to a CSV file.

The framework is implemented in Java and uses interfaces and abstract classes to extend user-defined components. `JMH` provides a strong foundation for constructing and executing profiling benchmarks whilst minimizing confounding runtime variables such as Java Virtual Machine (JVM) startup, Just-in-Time (JIT) warmup, Garbage Collection (GC) pauses.

The framework has extension for Octagon abstract domain, i.e., $Z = Octagon$. The implementation encodes Octagon constraints, which are constraints of the form: $\pm x \pm y \leq c$, where $x, y \in V$ where V is the set of variables and $c \in \mathbb{I}$, where $\mathbb{I}$ is one of $\mathbb{R}$, $\mathbb{Q}$, or $\mathbb{Z}$. Octagons are encoded as 2-dimensional array difference bounded matrix (DBM) [6] in the `OctagonDBM` class.

To extend operations over Octagons, users would provide extensions to `OctagonDBM`, overriding various operations with their implementation they wish to test. Furthermore, they would provide additional instances of `*Bench`, e.g., `JoinBench`. Similar to `JUnit` [4], the naming is conventional: `muRelBench` automatically includes classes containing the `Bench` suffix.

User extension beyond Octagons It is reasonably straightforward to extend `muRelBench` with additional abstract domains. A user must provide three additional classes: the abstract container type for the domain, e.g., `ZoneDBM` to add Zones [12]; a builder for the new abstract type; and finally, a state type which provides the different parameterization sets for `JMH`.

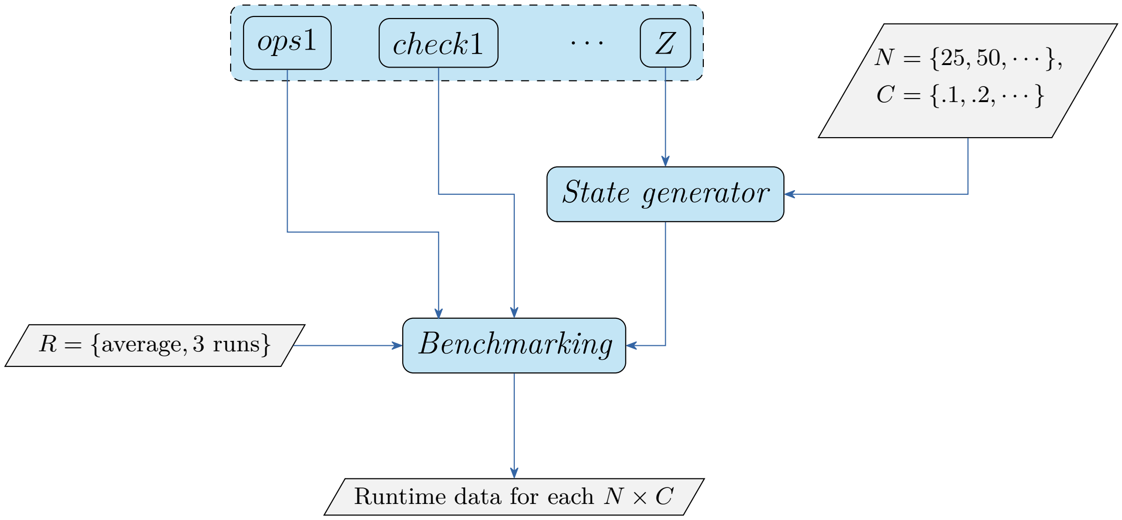

Fig. 1: Component diagram of `muRelBench`, specifying the framework's and user-defined components.

3 Octagons and Closure Operation Case Study

Benchmark Set Up We examine the benefits of `muRelBench` in a case study. The framework randomly generates Octagons, varying the "density" of relations between variables to create a continuum of synthetic instances of Octagons. This density progression roughly correlates to the different instances of Octagons from real programs. That is, early in analysis, variables have a tendency to have few relations as only few program statements are explored. In the middle of analysis, after exploring several assignment statements, variables become tightly coupled with one another. Finally, after several fixed point iterations and widening operations, islands of connectivity emerge [7,8,18]. Furthermore, we also vary the number of variables of the synthetic Octagons to account for different programs sizes.

For this case study, we generate Octagons with 25, 50, and 100 program variables, i.e., 50, 100, and 200 variables using the Octagon variable encoding [14]. For each size, we generate Octagons with $10\% - 90\%$ density, in 10% increments. The Cartesian product of these parameters results in 27 Octagon instances.

Using `JMH`, we default to 3 "warmup" iterations and 5 experimental iterations for each benchmark. Thus, for a single benchmark, the operation under test executes 216 times. However, we do provide options for the user to modify and otherwise specify their own desired warmup and experimental iterations, among other options available via `JMH`.

Case Study In this case study we chose to evaluate different closure algorithms for Octagon abstract domain. The closure represents a critical operation for static program analysis and abstract interpretation because it provides critical functions: normalization for equality comparisons for data-flow analysis (DFA) [1] and precision benefits for other domain operations such as LUB [13].

Canonicalizing or normalizing Octagon states is a necessary operation because an Octagonal bounded region can be represented by infinitely many different Octagons. The closure operation normalizes an Octagon by making explicit implicit edges and minimizing edge weights between variables within the Octagons. In the simplest case, this amounts to computing the all-pairs shortest-path problem for the directed, weighted graph used to represent the Octagon.

There exist several algorithms for computing the all-pairs-shortest-path problem for weighted-directed graphs such as Floyd-Warshall and Bellman-Ford algorithms [5]. While these algorithms are relatively simple and straightforward to implement, their cost can be excessive. Floyd-Warshall, for example has cubic time complexity, $\Theta(n^3)$, where n is the number of variables in the abstract Octagon state.

Chawdhary et al. [3] proposed an incremental closure algorithm for Octagons which uses code motion and hoisting to minimize the number of comparisons required to incrementally close an Octagon. Thus, they were able to reduce the incremental closure, a modified Floyd-Warshall, to $O(20n^2 - 4n)$.

Closure	**Program**	**Mean (ms)**	σ
Floyd-Warshall	Fibonacci	144	32.2
	Loop	46.8	3.1
Chawdhary	Fibonacci	117	5.1
	Loop	49.6	10.3

Table 1: Small programs used to demonstrate performance characteristics of using different closure algorithms.

Clearly, these two algorithms should have a different runtime growth with the increased number of variables. We first examined their result in the context of DFA on two small programs to see if any differences can be detected. Table 1 shows the results of the full-closure algorithm Floyd-Warshall and the Chawdhary et al's incremental closure. The data is averaged over five executions and includes the mean runtime for each along with their standard deviation. As the data shows that the results are not that conclusive since on Loop program Floyd-Warshall performed better while Chawdhary runs faster on Fibonacci. When we analyzed the properties of two program, we discovered that Fibonacci algorithm had the maximum of 6 variables with density of 72% and Loop program had two variables with no density, which is purely interval.

Plots in Fig. 2 show the results of the comparison of the two closure algorithm on benchmarks that `muRelBench` generates and runs. Each plot presents runtime data for different values of N while varying in density of connections between variables. Using this detailed data we can discern instances where the algorithm proposed by Chawdhary et al. outperforms full closure. Both algorithms are almost indistinguishable in the first two plots. For the plot with largest $N = 100$ we can see that, depending on the density value, one algorithm performs better than the other. For lower density states, Floyd-Warshall tends to perform better while for higher densities Chawdhary's incremental closure has better runtime data.

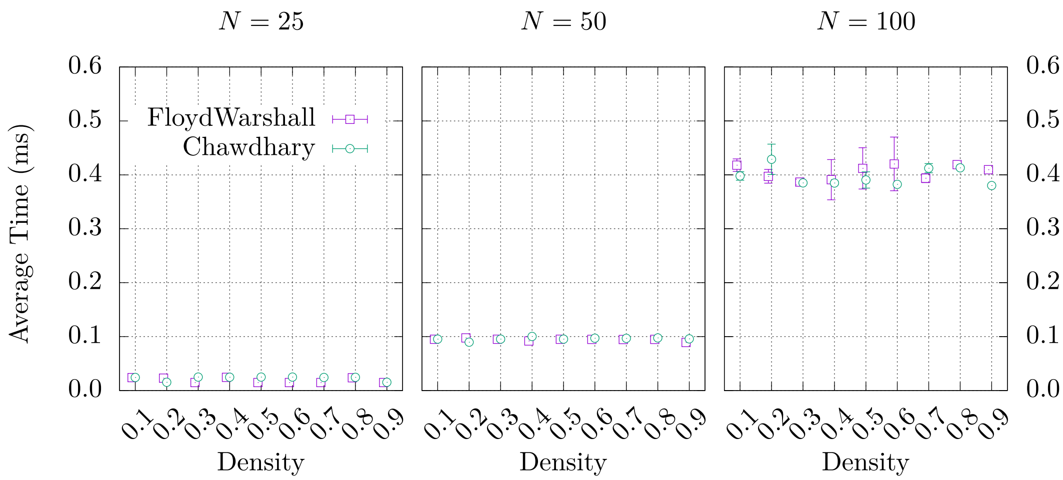

Fig. 2: Plots of microbenchmark results of closure operations, each subplot varies the number of variables, each sample varies the connectivity of program variables.

4 Conclusion and Future Work

In this paper, we present the abstract interpretation research community with `muRelBench` benchmarking framework that provides standardized and uniform support for comparing different operations within Zonotope abstract domains. When developing new algorithms or new abstract domains, a standard set of benchmarks and a framework to easily test them helps convince the community of their value.

Our framework of generated benchmarks invites many improvements and future work to make it better suited for the research community and software engineers at large. For example, we invite contributions of additional algorithms to be added to the suite, so others can use the results in their comparisons. Additionally, more parameters could provide a wider surface area of study for different Zonotope operations.

References

1. Ballou, K., Sherman, E.: Incremental transitive closure for zonal abstract domain. In: Deshmukh, J.V., Havelund, K., Perez, I. (eds.) NASA Formal Methods. pp. 800–808. Springer International Publishing, Cham (5 2022). `https://doi.org/10.1007/978-3-031-06773-0_43`, `http://dx.doi.org/10.1007/978-3-031-06773-0_43`
2. Brunner, R., Dyer, R., Paquin, M., Sherman, E.: Paclab: A program analysis collaboratory. In: Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering. ESEC/FSE '20, ACM (11 2020). `https://doi.org/10.1145/3368089.3417936`, `http://dx.doi.org/10.1145/3368089.3417936`
3. Chawdhary, A., Robbins, E., King, A.: Incrementally closing octagons. Formal Methods in System Design **54**(2), 232–277 (1 2018). `https://doi.org/10.1007/s10703-017-0314-7`, `http://dx.doi.org/10.1007/s10703-017-0314-7`
4. Contributors, M.: Junit (2024), `https://junit.org/`
5. Cormen, T., Leiserson, C., Rivest, R., Stein, C.: Introduction to Algorithms. Computer science, McGraw-Hill (2009). `https://doi.org/10.1.1.708.9446`, `https://books.google.com/books?id=aefUBQAAQBAJ`

6. Dill, D.L.: Timing assumptions and verification of finite-state concurrent systems. Lecture Notes in Computer Science pp. 197–212 (1990). https://doi.org/10.1007/3-540-52148-8_17, http://dx.doi.org/10.1007/3-540-52148-8_17
7. Gange, G., Ma, Z., Navas, J.A., Schachte, P., Søndergaard, H., Stuckey, P.J.: A fresh look at zones and octagons. ACM Transactions on Programming Languages and Systems **43**(3), 1–51 (9 2021). https://doi.org/10.1145/3457885, http://dx.doi.org/10.1145/3457885
8. Gange, G., Navas, J.A., Schachte, P., Søndergaard, H., Stuckey, P.J.: Exploiting sparsity in difference-bound matrices. Lecture Notes in Computer Science pp. 189–211 (2016). https://doi.org/10.1007/978-3-662-53413-7_10, http://dx.doi.org/10.1007/978-3-662-53413-7_10
9. Ghorbal, K., Goubault, E., Putot, S.: The zonotope abstract domain taylor1+. Lecture Notes in Computer Science pp. 627–633 (2009). https://doi.org/10.1007/978-3-642-02658-4_47, http://dx.doi.org/10.1007/978-3-642-02658-4_47
10. Jordan, M., Hayase, J., Dimakis, A., Oh, S.: Zonotope domains for lagrangian neural network verification. Advances in Neural Information Processing Systems **35**, 8400–8413 (2022)
11. Laaber, C., Leitner, P.: An evaluation of open-source software microbenchmark suites for continuous performance assessment. In: Proceedings of the 15th International Conference on Mining Software Repositories. ICSE '18, ACM (5 2018). https://doi.org/10.1145/3196398.3196407, http://dx.doi.org/10.1145/3196398.3196407
12. Miné, A.: A new numerical abstract domain based on difference-bound matrices. Lecture Notes in Computer Science pp. 155–172 (2001). https://doi.org/10.1007/3-540-44978-7_10, http://dx.doi.org/10.1007/3-540-44978-7_10
13. Miné, A.: Weakly Relational Numerical Abstract Domains (12 2004), https://pastel.archives-ouvertes.fr/tel-00136630
14. Miné, A.: The octagon abstract domain. Higher-Order and Symbolic Computation **19**(1), 31–100 (3 2006). https://doi.org/10.1007/s10990-006-8609-1, http://dx.doi.org/10.1007/s10990-006-8609-1
15. Oracle: jmh (2024), https://openjdk.org/projects/code-tools/jmh/
16. Oracle: openjdk/jmh (2024), https://github.com/openjdk/jmh
17. Schwarz, M., Seidl, H.: Octagons revisited. Lecture Notes in Computer Science p. 485–507 (2023). https://doi.org/10.1007/978-3-031-44245-2_21, http://dx.doi.org/10.1007/978-3-031-44245-2_21
18. Singh, G., Püschel, M., Vechev, M.: Making numerical program analysis fast. ACM SIGPLAN Notices **50**(6), 303–313 (8 2015). https://doi.org/10.1145/2813885.2738000, http://dx.doi.org/10.1145/2813885.2738000