Search-Guided Generation of Properties for Program Analyzers

^{1,2}Daniela Ferreiro

¹Universidad Politécnica de Madrid (UPM) Madrid, Spain ²IMDEA Software Institute, Madrid, Spain daniela.ferreiro@imdea.org

1 Introduction and Problem Description

Since the early years of logic programming, static program analysis has been a fundamental technique used to achieve several objectives ranging from sequential program optimization [36, 37, 25, 5] and parallelization [27, 3] to program verification [19, 31, 20, 32, 18]. However, static analyzers and verifiers are typically large and complex, which can make them prone to bugs. This can limit their applicability in real-life production compilers and development environments, since the tasks they are used in are typically critical and need reassurance about the soundness of the analysis results. At the same time, the validation of static analyzers is a challenging problem, with limited literature and few dedicated tools addressing this issue [4, 24, 38, 39, 9, 22, 15]. This is probably due to the fact that direct application of formal methods is not always straightforward with code that is so complex and large, even without considering the problem of having precise specifications to check against. This is a clear instance of the classic problem of who checks the checker. As a result, in current practice, extensive testing is the most extended and realistic validation technique, but this poses some significant challenges too. Testing separate components of the analyzer misses integration testing, and designing proper oracles for testing the complete tool is difficult.

A further complication lies in test data generation. Useful test data is often unavailable or too specialized, lacking the configurability needed to explore a wide range of scenarios. Even when data exists, it may fail to trigger corner case behaviors. In other words, achieving good coverage, especially for specific semantic properties, requires not only proper benchmarks but also diverse, high-quality inputs. Additionally, the generation of suitable properties and oracles to assess expected behavior is itself an open problem in many cases.

2 Research Objectives

In view of the above, our research objectives for the PhD are as follows:

- A first objective is the *development, improvement, and extensive benchmarking* of a practical technique for testing static analyzers. Our approach here is based on exploiting the particular characteristics of the program analyzers, verifiers, and assertion languages developed for (C)LP.
- A second objective involves the development of novel methods to generate relevant test cases, that effectively exercise a wide range of analysis scenarios. Our approach here is also based on exploiting the specific characteristics of (C)LP, and in particular using properties in assertions

© Daniela Ferreiro This work is licensed under the Creative Commons Attribution License. written as (C)LP programs as generators and flexible search rules to perform constraint-based automatic generation of test cases. This will involve *pushing the development of techniques for* supporting flexible search rules in general in (C)LP programs, which should also be useful beyond test case generation.

• A final objective is to prove practically all the proposed techniques in a realistic analyzer and compiler. To this end, we develop our work in the context of Ciao Prolog [16]. The Ciao programming environment includes CiaoPP, a large and complex abstract interpretation-based static analysis tool which faces the specific challenges that we are addressing. Recently, there has been some interesting work [33] aimed at verifying the partial correctness of the PLAI analysis algorithm (also referred to as "the top-down solver") that lies at the heart of CiaoPP using the Isabelle prover [30], but verification of the *actual implementation* remains a challenge. Like other "classic" analyzers, the CiaoPP formal framework has evolved for a long time, incorporating a large number of abstract domains, features, and techniques, adding up to over half a million lines of code. These components have, in turn, reached over the years different levels of maturity. While the essential parts, such as the fixpoint algorithms and the classic abstract domains, have been used routinely for a long time now, and it is unusual to find bugs, other parts are less developed and yet others are prototypes or even proofs of concept. We thus argue that the Ciao/CiaoPP system is thus a good subject for testing our techniques.

In the following sections, we explain further our progress so far in the thesis objectives and the remaining challenges towards these objectives.

Some of my prior related work: In my previous Bs and Ms work, I developed a tool for including real outputs from external tools in manuals and tutorials, so that they are always in sync with the actual tool being documented [1]. I applied these ideas to developing CiaoPP analysis and verification tutorials that evolve automatically with the system they describe, so that they are kept updated and without obsolete elements [1, 2]. I also integrated the Ciao analysis and verification tools in the browser within the Ciao playground, which is a powerful, in-browser tool for developing (C)LP code, notebooks, and presentations that contain runnable examples [26, 11], and that has been used extensively for teaching Prolog and (C)LP in general.

3 Current Progress and Results

3.1 The Checkification Algorithm

The first component of our work is the *checkification* algorithm and system [12]: a simple, automatic, technique for testing static analyzers. This work is based on preliminary work by I. Casso and other members of the group on the topic [8], but we have extended it with new ideas and techniques, as well as implemented it and fully benchmarked it. *Checkification* combines four basic components (all present in the Ciao system):

- A static analyzer (the PLAI static analyzer [28, 21, 14]),
- An assertion run-time checking framework [34, 35],
- A (random) test case generation framework [6],
- A unit-testing framework [23]



Figure 1: The Ciao assertion framework (CiaoPP's verification/testing architecture).

These elements are combined in a novel way that allows testing the static analyzer almost for free. Intuitively, the idea consists in checking, over a suite of benchmarks (or generated benchmarks), that the properties inferred statically are satisfied dynamically. The overall testing process (see Figure 1), for each benchmark, can be summarized as follows: first, the code is analyzed, obtaining the analysis results expressed as assertions interspersed within the original code. Then, the status of these assertions is switched into run-time checks, that will ensure that violations of those assertions are reported at execution time. Finally, test cases are generated and executed to exercise those run-time checks. Given that these assertions (the analyzer output) must cover all possible concrete executions (and assuming the correctness of our checking algorithm implementation), if any assertion violation is reported, assuming that the run-time checks are correct, it means that the assertion was incorrectly inferred by the analyzer, thus revealing an error in the analyzer itself. The error can, of course, sometimes also be in the run-time checks, but, typically, run-time checking is simpler than inference. This process is automatable, and, if it is repeated for an extensive and varied enough suite of benchmarks, it can be used to effectively validate (even if not fully verify) the analyzer or to discover new bugs. Furthermore, the implementation, when framed within a tool environment that follows the Ciao assertion model, is comparatively simple, at least conceptually.

The idea of checking at run time the properties or assertions inferred by the analysis for different program points is not new. For example, in [38], this technique was applied for checking a range of different aliasing analyses. However, these approaches require the development of tailored instrumentation or monitoring, and significant effort in their design and implementation. We argue that the testing approach is made more applicable, general, and scalable by the use of a unified assertion-based framework for static analysis and dynamic debugging, as the Ciao assertions model. As mentioned before, by developing the approach within such a framework, it can be implemented with many of the already existing algorithms and components in the system, in a very simple way.

Illustrative example. Let us illustrate the main idea of the approach with an example. Assume we have the following simple Prolog program, where we use an assertion to define the entry point for the analysis. The entry assertion indicates that the predicate is called with its second argument instantiated to a list, and the third a free variable.

```
1 :- entry prepend(_,+list,-).
2
3 prepend(X,Xs,Ys) :-
4 Ys=[X|Rest],
5 Rest=Xs.
```

Assume that we analyze it with a *simple modes* abstract domain that assigns to each variable in an abstract substitution one of the following abstract values:

- ground (the variable is ground),
- var (the variable is free),
- nonground (the variable is not ground),
- nonvar (the variable is not free),
- ngv (the variable is neither ground nor free), or
- any (nothing can be said about the variable).

Assume also that the analysis is incorrect because it does not consider sharing (aliasing) between variables, so when updating the abstract substitution after the Rest=Xs literal, the abstract value for Ys is not modified at all. The result of the analysis will be represented as a new source file with interspersed assertions, as shown in Fig. 2 (lines 3-5, 8, 10, and 12). Note that the correct result, if the analysis considered aliasing, would be that there is no groundness information for Ys at the end of the clause (line 12), since there is none for X or Xs at the beginning either. Ys could only be inferred to be nonvar, but instead is incorrectly inferred to be nonground too (line 10). Normally any/1 properties (i.e., top, or unknown) would not actually be included in the analysis output for conciseness, but are included in Fig. 2 for clarity.

The objective of our approach is to check dynamically the validity of these true assertions from the analyzer, that in this case contains an error. The insight is that, thanks to the different capabilities of the Ciao model this can be achieved by (1) *turning the status of the true assertions produced by the analyzer into check*, as shown in Fig. 3. This would normally not make any sense since these true assertions have been proved by the analyzer. But that is exactly what we want to check, i.e., whether the information inferred is incorrect. To do this, (2) we run the transformed program (Fig. 3) again through CiaoPP (Fig. 1) but without performing any analysis.

In that case, the check literals (stemming from the true literals of the previous run) will not be simplified in the comparator (since there is no abstract information to compare against) and instead will

```
:- entry prepend(_,+list,-).
1
2
     :- true pred prepend(X,Xs,Ys)
3
         : (any(X), nonvar(Xs), var(Ys))
4
         => (any(X), nonvar(Xs), nonground(Ys), nonvar(Ys)).
5
6
    prepend(X,Xs,Ys) :-
7
         true(any(X), nonvar(Xs), var(Ys), var(Rest)),
8
         Ys=[X | Rest],
9
         true(any(X), nonvar(Xs), nonground(Ys), nonvar(Ys), var(Rest)),
10
         Rest=Xs.
11
         true(any(X), nonvar(Xs), nonground(Ys), nonvar(Ys), nonvar(Rest)).
12
```

Figure 2: An incorrect simple mode analysis.

```
:- entry prepend(_,+list,-).
1
2
3
     :- check pred prepend(X,Xs,Ys)
         : (any(X), nonvar(Xs), var(Ys))
4
5
         => (any(X), nonvar(Xs), nonground(Ys), nonvar(Ys)).
6
7
    prepend(X,Xs,Ys) :-
         check(any(X), nonvar(Xs), var(Ys), var(Rest)),
8
9
         Ys=[X|Rest].
         check(any(X), nonvar(Xs), nonground(Ys), nonvar(Ys), var(Rest)),
10
11
         Rest=Xs.
         check(any(X), nonvar(Xs), nonground(Ys), nonvar(Ys), nonvar(Rest)).
12
```

Figure 3: The instrumented program.

be converted directly to run-time tests. In other words, the check(*Goal*) literals will be expanded and compiled to code that, every time that this program point is reached, in every execution, will check dynamically if the property (or properties) within the check literal (i.e., those in *Goal*) succeed, and an error message will be emitted if they do not.

The only missing step to complete the automation of the approach is to (3) run prepend/3 on a set of test cases. These may, in general, already be available as test assertions in the program or, alternatively, the random test case generator can be used to generate them. E.g., for prepend/3 the test generation framework will ensure that instances of the goal prepend(X, Xs, Ys) are generated, where Xs is constrained to be a list, and Ys remains a free variable. However, X and the elements of Xs will otherwise be instantiated to random terms. In this example, as soon as a test case is generated where both X and all elements in Xs are ground, the program will report a run-time error in the check in line 12, letting us know that the third program point, and thus the analysis, is incorrect. The same procedure can be followed to debug different analyses with different benchmarks. If the execution of any test case reports a run-time error for one assertion, it will mean that the assertion was not correct and the analyzer computed an incorrect over-approximation of the semantics of the program. Alternatively, if this experiment, which can be automated easily, is run for an extensive suite of benchmarks without errors, we can gain more confidence that our analysis implementation is correct, even if perhaps imprecise (although of course, we cannot have actual correctness in general by testing).

We have carried out a full implementation of the approach and completed a substantial experimental evaluation [12]. The results show that checkification can effectively discover and locate interesting, non-trivial and previously undetected bugs, with reasonable overhead, not only in the less-developed parts of the system, but also in corner cases of the more mature components, such as the handling of built-ins, run-time checking instrumentation, etc. The approach has also proven useful for detecting issues in auxiliary stages of analysis and verification, including assertion simplification, pretty printing, abstract program optimizations and transformations, etc. The results and more details about the algorithm are provided in [12].

3.2 Search Rules and Generation of Test Cases from Properties

As mentioned before, a second challenge that we are addressing is that the test case generation process that is necessary during the validation of static analyzers using checkification, and, in fact, while testing logic programs in general, becomes progressively difficult when dealing with complex data structures and programs whose correctness depends on complex, hard-to-test conditions. Prolog provides a unique

```
:- search_rule(tree/1,bf).
1
2
3
   :- prop tree/1.
4
  tree(empty).
5
  tree(t(L,N,R)) :-
       tree(L),
6
7
       gen([sr(rnd)], (nnegint(X), X =< 15)),</pre>
8
       tree(R).
9
  % Check if tree T is sorted
10
   sorted_tree(T) :- ...
11
12
  % tsum(T,N) : Constraints the sum of all node values in tree T to equal N
13
  tsum(T,N) :- ...
14
15
  :- pred insert(X,T0,T1) : tree(T0) => tree(T1) + gen([sorted_tree(T0), tsum(T0,10)]).
16
  insert(X,T0,T1) :- ...
17
18
  ... % rest of the implementation of module predicates
19
```

advantage in this context, as its declarative nature allows the properties to be proved to be expressed as predicates, which can then be leveraged to generate test cases [7, 13, 12, 10, 29].

Prolog systems traditionally employ depth-first search as their execution strategy. While this choice is well-justified for efficiency reasons, and generally accepted, it is well known that this can lead to incompleteness when evaluating programs over infinite search spaces. This second thesis objective motivates us to revisit the role of the search rule in Prolog and explore the challenges involved in running Prolog predicates with alternative search strategies in the real world. The hypothesis is that flexible and customizable control mechanisms can be very useful in uncovering edge cases in testing, particularly when generating data structures with strong invariants. At the same time, our objective is to *push the development of techniques for supporting flexible search rules in general in (C)LP programs*, to be useful beyond test case generation.

Writing Prolog code that implements a search with a particular strategy from scratch is not particularly difficult and it is also not difficult to run Prolog predicates using other search rules using variations of the standard meta-interpreter. However, our challenge is to be able to run standard Prolog predicates with different search strategies while maintaining compatibility with modules, built-ins, and other libraries and features. Straightforward implementation approaches typically limit the use of alternative search rules to a subset of Prolog and thus do not support the full expressiveness of the language.

In the context of test case generation, our work leverages assertion preconditions as generators: since these preconditions are conjunctions of literals, the corresponding predicates can be used to systematically produce valid inputs. The key innovation lies in executing standard predicates under non-standard search rules, enabling either fully automatic or user-guided generation. Figure 3.2 demonstrates this approach with a binary tree library.

The tree/1 property describes the shape of trees, allowing either an empty node or a compound term with left and right subtrees. The use of gen([sr(rnd)], (nnegint(X), X =< 15)) in the recursive clause shows how value generation can be customized: in this case, generating non-negative integers bounded by 15. Moreover, the declaration :- search_rule(tree/1, bf) specifies that breadth-first search should be used to generate trees, which helps avoid unbalanced growth and improves coverage by generating smaller and more diverse trees earlier. In the insert/3 specification, the generator is further guided by auxiliary properties such as sorted_tree/1 and tsum/2, which restrict the structure and content of the input trees.

Our thesis is that flexible search rules can significantly enhance test case generation, enabling users to customize and implement their own generators, including those producing data structures with complex invariants.

4 Conclusions and Future Work

Our starting point has been *checkification*, an automatic method for testing static analysis tools by checking that the properties inferred statically are satisfied dynamically. We have shown how checkification can be implemented effectively in practice, improving several aspects of the approach, and performed a full benchmarking of the technique showing its usefulness in practice for discovering and locating interesting, non-trivial and previously undetected bugs, with reasonable overhead.

During this first phase of the work, we have identified test case generation as an important remaining challenge, and as a result, as our second thesis objective, we are developing methods for specifying search procedures in a concise and elegant way for realistic Prolog programs. Our work here allows users to switch search rules at any point in the program, and to combine simple primitives to build useful search heuristics for generating complex data structures for test cases. Future work in the thesis involves improving the generation process by supporting more configurable search parameters, including termination conditions, solution limits, and other customizable aspects. We will apply these techniques to the generation of programs from grammars for creating benchmarks which can serve as sophisticated automatic inputs to further automate the *checkification* algorithm.

While testing approaches are obviously ultimately insufficient for proving the correctness of analyzers, and thus it is clearly worthwhile to also pursue the avenue of code verification, we believe that the approaches addressed in the thesis will together offer a practical and effective technique for detecting errors in large and complex analysis tools, as well as in logic programs in general. In addition, we believe that the availability of sophisticated and easy to use search rules in Prolog can also be useful in both teaching the language [17] and in practical applications.

References

- [1] Daniela Ferreiro de Aguiar (2021): *Automatic Analysis of Code Examples*. Master's thesis, Universidad Politécnica de Madrid, ETSIInf, E-28660, Boadilla del Monte, Madrid, Spain. BSc Thesis.
- [2] Daniela Ferreiro de Aguiar (2023): *A System for generating interactive tutorials for CiaoPP*. Master's thesis, Universidad Politécnica de Madrid, ETSIInf, E-28660, Boadilla del Monte, Madrid, Spain. MSc Thesis.
- [3] F. Bueno, M. García de la Banda & M.V. Hermenegildo (1999): *Effectiveness of Abstract Interpretation in Automatic Parallelization: A Case Study in Logic Programming. ACM TOPLAS* 21(2), pp. 189–238.
- [4] Alexandra Bugariu, Valentin Wüstholz, Maria Christakis & Peter Müller (2018): Automatically Testing Implementations of Numerical Abstract Domains. In: Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering, ASE 2018, Association for Computing Machinery, New York, NY, USA, p. 768–778, doi:10.1145/3238147.3240464. Available at https://doi.org/10.1145/3238147.3240464.
- [5] M. Carro, J. Morales, H.L. Muller, G. Puebla & M.V. Hermenegildo (2006): *High-Level Languages for Small Devices: A Case Study*. In Krisztian Flautner & Taewhan Kim, editors: Compilers, Architecture, and Synthesis for Embedded Systems, ACM Press / Sheridan, pp. 271–281.
- [6] I. Casso, J. F. Morales, P. Lopez-Garcia & M.V. Hermenegildo (2020): An Integrated Approach to Assertion-Based Random Testing in Prolog. In Maurizio Gabbrielli, editor: Post-Proceedings of the 29th International

Symposium on Logic-based Program Synthesis and Transformation (LOPSTR'19), LNCS 12042, Springer-Verlag, pp. 159–176, doi:10.1007/978-3-030-45260-5_10.

- [7] I. Casso, J. F. Morales, P. Lopez-Garcia & M.V. Hermenegildo (2020): An Integrated Approach to Assertion-Based Random Testing in Prolog. In Maurizio Gabbrielli, editor: 29th International Symposium on Logicbased Program Synthesis and Transformation (LOPSTR'19), LNCS 12042, Springer-Verlag, pp. 159–176, doi:10.1007/978-3-030-45260-5_10.
- [8] Ignacio Casso, José F. Morales, Pedro López-García & Manuel V. Hermenegildo (2021): *Testing Your (Static Analysis) Truths*. In Maribel Fernández, editor: Logic-Based Program Synthesis and Transformation 30th International Symposium, Post-Proceedings, Lecture Notes in Computer Science 12561, Springer, pp. 271– 292, doi:10.1007/978-3-030-68446-4_14.
- [9] Pascal Cuoq, Benjamin Monate, Anne Pacalet, Virgile Prevosto, John Regehr, Boris Yakobowski & Xuejun Yang (2012): *Testing Static Analyzers with Randomly Generated Programs*. In Alwyn E. Goodloe & Suzette Person, editors: NASA Formal Methods, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 120–125.
- [10] Richard Denney (1991): Test-Case Generation from Prolog-Based Specifications. IEEE Software 8(2), pp. 49–57.
- [11] D. Ferreiro, J.F. Morales, S. Abreu & M.V. Hermenegildo (2023): Demonstrating (Hybrid) Active Logic Documents and the Ciao Prolog Playground, and an Application to Verification Tutorials. In: Technical Communications of the 39th International Conference on Logic Programming (ICLP 2023), Electronic Proceedings in Theoretical Computer Science (EPTCS) 385, Open Publishing Association (OPA), pp. 324–330. Available at https://cliplab.org/papers/hald-demo-iclp-tc.pdf. See also associated poster at https://cliplab.org/papers/hald-poster-iclp.pdf.
- [12] Daniela Ferreiro, Ignacio Casso, Pedro Lopez-Garcia, José F. Morales & Manuel V. Hermenegildo (2025): Checkification: A Practical Approach for Testing Static Analysis Truths. Theory and Practice of Logic Programming. Available at https://arxiv.org/abs/2501.12093.
- [13] Sophie Fortz, Fred Mesnard, Étienne Payet, Gilles Perrouin, Wim Vanhoof & Germán Vidal (2020): An SMT-Based Concolic Testing Tool for Logic Programs. In Keisuke Nakano & Konstantinos Sagonas, editors: Functional and Logic Programming - 15th International Symposium, FLOPS 2020, Akita, Japan, September 14-16, 2020, Proceedings, Lecture Notes in Computer Science 12073, Springer, pp. 215–219, doi:10.1007/978-3-030-59025-3_13.
- [14] I. Garcia-Contreras, J.F. Morales & M.V. Hermenegildo (2020): Incremental Analysis of Logic Programs with Assertions and Open Predicates. In: Proceedings of the 29th International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR'19), LNCS 12042, Springer, pp. 36–56, doi:10.1007/978-3-030-45260-5_3.
- [15] Weigang He, Peng Di, Mengli Ming, Chengyu Zhang, Ting Su, Shijie Li & Yulei Sui (2024): Finding and Understanding Defects in Static Analyzers by Constructing Automated Oracles. In: ACM International Conference on the Foundations of Software Engineering, doi:10.1145/3660781.
- [16] M. V. Hermenegildo, F. Bueno, M. Carro, P. Lopez-Garcia, E. Mera, J.F. Morales & G. Puebla (2012): An Overview of Ciao and its Design Philosophy. Theory and Practice of Logic Programming 12(1-2), pp. 219-252, doi:10.1017/S1471068411000457. Available at https://arxiv.org/abs/1102.5497.
- [17] Manuel V. Hermenegildo, José F. Morales & Pedro Lopez-Garcia (2024): *Teaching Pure LP with Prolog and a Fair Search Rule*. In: *Proceedings of the 40th ICLP Workshops*, 3799, CEUR-WS.org. Available at https://ceur-ws.org/Vol-3799/paper2PEG2.0.pdf.
- [18] M.V. Hermenegildo, J.F. Morales, P. Lopez-Garcia & M. Carro (2023): Types, modes and so much more the Prolog way. In David S. Warren, Veronica Dahl, Thomas Eiter, Manuel V. Hermenegildo, Robert Kowalski & Francesca Rossi, editors: Prolog - The Next 50 Years, chapter 2, LNCS 13900, Springer, pp. 23–37. Available at https://cliplab.org/papers/AssertionsAndOther-PrologBook.pdf.

- [19] M.V. Hermenegildo, G. Puebla & F. Bueno (1999): Using Global Analysis, Partial Specifications, and an Extensible Assertion Language for Program Validation and Debugging. In: The Logic Programming Paradigm: a 25–Year Perspective, Springer-Verlag, pp. 161–192.
- [20] M.V. Hermenegildo, G. Puebla, F. Bueno & P. Lopez Garcia (2005): Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor). Science of Computer Programming 58(1–2), pp. 115–140.
- [21] M.V. Hermenegildo, G. Puebla, K. Marriott & P. Stuckey (2000): Incremental Analysis of Constraint Logic Programs. ACM Transactions on Programming Languages and Systems 22(2), pp. 187–223.
- [22] Christian Klinger, Maria Christakis & Valentin Wüstholz (2019): Differentially Testing Soundness and Precision of Program Analyzers. In: Proceedings of the 28th ACM SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2019, Association for Computing Machinery, New York, NY, USA, p. 239–250, doi:10.1145/3293882.3330553. Available at https://doi.org/10.1145/3293882.3330553.
- [23] E. Mera, P. Lopez-Garcia & M.V. Hermenegildo (2009): Integrating Software Testing and Run-Time Checking in an Assertion Verification Framework. In: 25th Int'l. Conference on Logic Programming (ICLP'09), LNCS 5649, Springer-Verlag, pp. 281–295.
- [24] Jan Midtgaard & Anders Møller (2017): QuickChecking Static Analysis Properties. Softw. Test., Verif. Reliab. 27(6), doi:10.1002/stvr.1640. Available at https://doi.org/10.1002/stvr.1640.
- [25] J. Morales, M. Carro & M.V. Hermenegildo (2004): Improving the Compilation of Prolog to C Using Moded Types and Determinism Information. In: Proceedings of the Sixth International Symposium on Practical Aspects of Declarative Languages, Lecture Notes in Computer Science 3057, Springer-Verlag, Heidelberg, Germany, pp. 86–103.
- [26] J.F. Morales, S. Abreu, D. Ferreiro & M.V. Hermenegildo (2022): *Teaching Prolog with Active Logic Doc-uments*. Technical Report CLIP-1/2022.0, Technical University of Madrid (UPM) and IMDEA Software Institute.
- [27] K. Muthukumar & M.V. Hermenegildo (1989): Determination of Variable Dependence Information at Compile-Time Through Abstract Interpretation. In: 1989 North American Conference on Logic Programming, MIT Press, pp. 166–189.
- [28] K. Muthukumar & M.V. Hermenegildo (1992): Compile-time Derivation of Variable Dependency Using Abstract Interpretation. Journal of Logic Programming 13(2/3), pp. 315–347.
- [29] N. Mweze & W. Vanhoof (July 2006): Automatic Generation of Test Inputs for Mercury Programs. In: Pre-proceedings of LOPSTR 2006. Extended abstract.
- [30] L.C. Paulson (1990): Isabelle: The Next 700 Theorem Provers. In P. Odifreddi, editor: Logic and Computer Science, Academic Press, pp. 361–386.
- [31] G. Puebla, F. Bueno & M.V. Hermenegildo (2000): A Generic Preprocessor for Program Validation and Debugging. In: Analysis and Visualization Tools for Constraint Programming, LNCS 1870, Springer-Verlag, pp. 63–107.
- [32] M.A. Sanchez-Ordaz, I. Garcia-Contreras, V. Perez-Carrasco, J. F. Morales, P. Lopez-Garcia & M.V. Hermenegildo (2021): VeriFly: On-the-fly Assertion Checking via Incrementality. Theory and Practice of Logic Programming 21(6), pp. 768–784.
- [33] Yannick Stade, Sarah Tilscher & Helmut Seidl (2024): Partial Correctness of the Top-Down Solver. Archive of Formal Proofs. https://isa-afp.org/entries/Top_Down_Solver.html, Formal proof development.
- [34] N. Stulova, J. F. Morales & M.V. Hermenegildo (2015): Practical Run-time Checking via Unobtrusive Property Caching. Theory and Practice of Logic Programming, 31st Int'l. Conference on Logic Programming (ICLP'15) Special Issue 15(04-05), pp. 726–741, doi:10.1017/S1471068415000344. Https://arxiv.org/abs/1507.05986.

- [35] N. Stulova, J. F. Morales & M.V. Hermenegildo (2016): Reducing the Overhead of Assertion Run-time Checks via Static Analysis. In: 18th Int'l. ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP'16), ACM Press, pp. 90–103.
- [36] P. Van Roy (1994): 1983-1993: The Wonder Years of Sequential Prolog Implementation. Journal of Logic Programming 19/20, pp. 385–441.
- [37] P. Van Roy & A.M. Despain (1992): *High-Performance Logic Programming with the Aquarius Prolog Compiler. IEEE Computer Magazine*, pp. 54–68.
- [38] Jingyue Wu, Gang Hu, Yang Tang & Junfeng Yang (2013): Effective Dynamic Detection of Alias Analysis Errors. In: Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2013, Association for Computing Machinery, New York, NY, USA, p. 279–289, doi:10.1145/2491411.2491439. Available at https://doi.org/10.1145/2491411.2491439.
- [39] Chengyu Zhang, Ting Su, Yichen Yan, Fuyuan Zhang, Geguang Pu & Zhendong Su (2019): *Finding and understanding bugs in software model checkers*. In: Proceedings of the 13th Joint Meeting of the 18th European Software Engineering Conference and the 27th Symposium on the Foundations of Software Engineering, pp. 763–773, doi:10.1145/3338906.3338932.