

SOUFFLÉ: On Synthesis of Program Analyzers

Herbert Jordan¹, Bernhard Scholz², and Pavle Subotic³(✉)

¹ University of Innsbruck, Innsbruck, Austria

² University of Sydney, Sydney, Australia

³ University College London, London, UK

psubotic@gmail.com

Abstract. SOUFFLÉ is an open source programming framework that performs static program analysis expressed in Datalog on very large code bases, including points-to analysis on OpenJDK7 (1.4M program variables, 350K objects, 160K methods) in under a minute. SOUFFLÉ is being successfully used for Java security analyses at Oracle Labs due to (1) its high-performance, (2) support for rapid program analysis development, and (3) customizability. SOUFFLÉ incorporates the highly flexible Datalog-based program analysis paradigm while exhibiting performance results that are on-par with manually developed state-of-the-art tools. In this tool paper, we introduce the SOUFFLÉ architecture, usage and demonstrate its applicability for large-scale code analysis on the OpenJDK7 library as a use case.

1 Introduction

Among the reasons for the slow industrial adoption of static program analysis is the lack of sufficient customizability and scalability in tools. Recently, the use of Datalog-like languages, has had a resurgence in several computer science communities [9], particularly, in the area of program analysis [2–4, 12, 16, 18] where tools such as μ Z [10], LogicBlox [11] and bddb [18] have shown great promise. In these tools, Datalog acts as a domain specific language to express custom program analyses concisely, reducing the complexity of developing program analyzers. The drawback of this approach is that program analyses specified in Datalog typically experience reduced performance compared to manually implemented tools. A notable reason for this decrease in performance appears to be the “one size fits all” approach of evaluating Datalog programs, i.e., Datalog engines generally lack the ability to specialize their evaluation process for a given instance of a program analysis specification.

To close the performance gap, we have developed a tool called SOUFFLÉ that overcomes the performance limitations of standard Datalog evaluation by performing an efficient synthesis of Datalog specifications to executable C++ programs. As a result, SOUFFLÉ is able to perform analyses on-par with state-of-the-art manual tools while retaining the advantages of employing a domain specific

Parts of this research was conducted while visiting Oracle Labs, Australia as assistants and visiting professor.




Fig. 1. Comparison: standard Datalog evaluation versus the architecture of SOUFFLÉ

language for expressing static program analyses. For example, [6] reports the ground-breaking capability of obtaining points-to analysis results for the Open-JDK library in under a minute. With the same dataset, SOUFFLÉ can obtain a similar performance (35s) using a general purpose analysis infrastructure on a multi-core commodity desktop system.

In this tool paper, we give an overview of the SOUFFLÉ framework; notably its architecture, optimizations and expected performance on very large code bases. We conclude with a summary of on-going developments of the SOUFFLÉ infrastructure.

2 How It Works

A Datalog program [1] consists of an extensional database, which is defined by facts, and an intensional database, which is defined by rules. In a setup for static program analysis, the extensional database represents an input program in relational form. The relational representation of an input program is obtained from an extractor [15] describing the relevant semantics of the input program for a given program analysis. The intensional database represents the program analysis specification phrased as Horn clause formulae over finite domains. Figure 1(a) illustrates the workflow for static program analysis in Datalog. The query result of the Datalog execution represents the actual result of the program analysis. While standard schemes for evaluating Datalog are generally optimized for reducing the amount of redundant computation, e.g., the conventional, interpreter-based semi-naïve evaluation scheme [1] as shown in Fig. 1(b), they lack the ability to specialize their evaluation for a given program analysis specification instance.

SOUFFLÉ takes a different approach: Instead of evaluating a Datalog program on-the-fly, we treat a Datalog program as a specification that is synthesized to a C++ program. The C++ program is compiled, and executed with the extensional database (i.e. facts) as an input. Essentially, the generated executable becomes an analyzer in its own right. Figure 1(c) depicts our translation scheme, where the Datalog specification is first parsed and semantically checked. The input specification is then translated internally to an imperative Relational Algebra Machine (RAM) program. The RAM program is further translated to a




Fig. 2. Specialization Hierarchy and First Futamura Projection of Semi-Naïve Evaluation

C++ program with OpenMP annotations for parallel execution and C++ template based meta-programming elements. In the last stage, an OpenMP/C++ compiler translates the generated code to a highly optimized, parallel program.


2.1 A Hierarchy of Specializations

To achieve a synthesis of Datalog specifications to C++, we follow a staged specialization hierarchy as depicted in Fig. 2(a). At each stage, a specialization step, as characterized by Futamura projections [8], is applied. The foundation is provided by an abstract transformation Mix that, given an interpreter Int and a source program Source , yields a specialized program amalgamating the interpreter and the source program. The specialized program performs the same computation as the source program (executed by the interpreter) – yet more efficiently. In Fig. 2(b), the semantic equivalence is shown between evaluation under an interpreter Int and the program produced by the Mix transformation [8]. What is of particular interest, is that at each specialization phase, information is revealed that enables opportunities for code optimizations that were not possible at earlier stages. As a consequence, the binary code produced by our specialization hierarchy is on-par in terms of run-time and memory usage with state-of-the-art hand crafted code.

The first specialization $P_{\text{RAM}} = \text{Mix}(\text{Int}_{d1}, \text{Idb})$ sees the Semi-Naïve Evaluation [1] as the interpreter, Int_{d1} . It is specialized with the intensional database, Idb , corresponding to the analysis specification. As a result, we receive a relational algebra machine program P_{RAM} that expresses the computation of the specified analysis as a series of fix-point computation steps over relational algebra operations. From a high-level viewpoint, the specialization of the Semi-Naïve evaluation is a translation of a declarative Datalog program to an imperative relational algebra program.

The next application of Futamura's projection is performed on the RAM program, i.e., $P_{\text{C++}\langle \rangle} = \text{Mix}(\text{Int}_{\text{RAM}}, P_{\text{RAM}})$, that has been generated by the first stage and the RAM interpreter Int_{RAM} . The conducted specializations target the efficient structuring of loop-based join operations and the identification of optimal index support, in order to reduce the worst-case runtime complexity

```
void m(int i, int j){
  s: while (i < j){
11:  protect();
12:  i++;
  }
13:  vulnerable();
}
```



```
// Declarations
E(s:Node,d:Node) input
P(node:Node) input
I(node:Node) output
// Analysis
I("s").
I(y) :- I(x), E(x,y), !P(y).
```

Fig. 3. Java-like input program, a graphical representation of its control-flow, and security specification in SOUFFLÉ

of the RAM program. However, index management is expensive and a minimal number of indices is desirable. In SOUFFLÉ we employ a novel, optimal and polynomial-time algorithm that is inspired by Dilworth’s theorem [7] to compute only necessary indices. The idea of the algorithm is to compute partitions of chains in a lattice of indices. From each chain, a maximum index is computed that subsumes all other indices in the chain. This optimization results in a large run-time improvement in the resulting analyzer. After specializing the relational algebra program, C++ code that makes extensive use of templates is generated.

The final specialization step, $P_{C++} = \text{Mix}(\text{Int}_{<>}, P_{C++<>})$, is performed while compiling the generated C++ program. This Futamura projection is implemented using template-based meta-programming techniques [17]. With meta-programming techniques, data structures and algorithms are specialized by static information, thereby hoisting computations from run-time to compile-time. E.g., data structure interfaces are realized in form of C++ concepts rather than polymorphic C++ base classes to eliminate virtual-call dispatches and run-time type checks. The generated data structures are highly specialized towards the use of the corresponding relations in the input program. We employ efficient parallel variations of B-trees and Tries, with customized data element, node, and iterator types. Additionally, primary and secondary index support is provided for efficient operations on the represented relation in the program. For example, one of the most time-consuming operations with the use of indices is the comparison of two tuples. For this purpose we instantiate specialized versions of templated lexicographical order functions in order to removing unnecessary control-flow and memory access overhead from the analysis run-time.

Our staged-translation approach using a specialization hierarchy coupled with standard Datalog optimizations and specialized relational data structures allows SOUFFLÉ to analyze very large code bases, previously considered to be impractical for Datalog-based engines. The generated C++ code is packaged in form of header files for a smooth integration with host applications.

2.2 An Example of the Specialization Process

Figure 3 illustrates a simple security analysis for an example assuming that there is a low and high security state in a program. The invocation of a security sensitive method `vulnerable` is permitted only in the high-security state. A call to the method `protect` transfers the security state from low to high if permitted.

<pre> 1 I = {s}; ΔI = I 2 while ΔI ≠ ∅ 3 I' = ∅ 4 for u ∈ ΔI do 5 for (x, y) ∈ E do 6 if (u = x ∧ y ∉ P ∧ y ∉ I) 7 I' = I' ∪ {y} 8 I = I ∪ I' 9 ΔI = I'</pre>	<pre> 1 I = {s}; ΔI = I 2 while ΔI ≠ ∅ 3 I' = ∅ 4 for u ∈ ΔI do 5 for (., y) ∈ E(u, .) do 6 if (y ∉ P ∧ y ∉ I) 7 I' = I' ∪ {y} 8 I = I ∪ I' 9 ΔI = I'</pre>
(a) Futamura Projection	(b) Partial Evaluation of RAM Program

Fig. 4. Running example

control-flow graph and m is the number of edges in the control-flow graph. To improve the performance of the program, we specialize the loop traversal of the loop in line 5 by employing an index. The index filters out all pairs in the edge relation whose source is not node u , i.e., all the edges are selected which emanate of node u denoted by the set $E(u, _)$. This specialization requires an index on relation E , yet significantly reduces the runtime complexity. Typical analyses result in potentially hundreds of indices, making index management expensive if performed naively. We therefore employ an optimal, minimal index selection technique based on Dilworth’s theorem [7] to select only necessary indices since an index may subsume several indices. Suppose we had another access to relation E on both u and v attributes, i.e., $E(u, v)$. A naive implementation would be to have two indices defined by the lexicographical orders u and $u < v$, however, the minimal solution would be to have only one index, namely, $u < v$ as it subsumes the index with only u . Some information to our solution to this combinatorial problem can be found in [13].

To implement indices from the previous step, we employ templated B-Trees that require a comparison function for two tuples in the relation. The comparison function is implemented as a lexicographical order in the form of a template as sketched below,

```

template<int...> struct Comperator;
template<int i, int ... tail > struct Comperator<i,tail...> {
    static bool cmp(const tuple& a, const tuple& b){
        return a[i] < b[i] || (a[i] == b[i] && Comperator<tail...>::cmp(a,b));
    }
};
template<> struct Comperator<> {
    static bool cmp(const tuple&, const tuple&){ return true; }
};
```

The variadic template for the struct `Comperator` is parametrized by the columns in order. E.g., the call `Comperator<2,0>::cmp(a,b)` compares the tuples `a` and `b` by checking whether the third element of `a` is less than the third element of `b`. If the comparison results in a tie, the first elements of both tuples are compared to determine the order between the two tuples `a` and `b`. The operator is defined recursively: the base case is given by the `struct Comperator<>` considering

every tuple equal, and the inductive case by `struct Comperator<i,tail...>`, comparing the i -th components and, if equal, delegating the comparison to `Comperator<tail...>`. The expansion of the template for a given instance such as `Comperator<2,0>` is performed at compile-time and delivers, in combination with function inlining, significant performance gains for index construction and retrieval. Without applying meta-programming techniques that rely on program specializations, i.e., pushing computations from runtime to compile-time, these performance gains would not be achievable.

3 Case Study: OpenJDK7

In this section we present our experience using SOUFFLÉ as a Java security analysis tool on the Java Development Kit (JDK). We point the reader to [5] for information on the Java vulnerabilities work at Oracle. For more detailed performance data on the techniques used in SOUFFLÉ, we refer the reader to [13].

In Table 1 we present three types of analyses performed on the OpenJDK7-b147. Due to the sheer size of OpenJDK7 (1.4M variables, 350K heap objects, 160K methods, 590K invocations and 17K types) such analyses are typically regarded as either impractical for most tools or at the very least, extremely challenging. The CI column refers to a *context-insensitive* points-to analysis and the CS refers to a *context-sensitive* points-to analysis. Points-to analysis is the main building block of most security analyses performed and are typically dominating the overall execution time. The last column, Security, refers to a large, composite security analysis similar to the *caller sensitive method* analysis in [5].

For our evaluation, we compare the performance of bddbdb [18], Z3’s Datalog extension μZ [10], an SQLite based Datalog engine [14], and a 8-core parallel version of SOUFFLÉ. Each analysis has been ported to the respective Datalog variation of the evaluated tools. The resulting specifications typically comprise a few thousand lines of code. For the SOUFFLÉ based specifications, SOUFFLÉ’s module system has been utilized to facilitate the reuse of code among the three analysis, reducing the necessary development effort.

Our experiments reveal the limited capability of pre-existing Datalog-based tools when analyzing very large code bases. The CI analysis represents a very simple points-to analysis that does not construct the call-graph on the fly. A

Table 1. Comparison of Datalog evaluation tools for analyses on the OpenJDK7 b147 library, executing on an 8 core Intel Xeon E5-2690 v2 @ 3.0 GHz server system. DNF = Did Not Finish within 18 h.

Tool	CI		CS		Security	
	Δt [hh:mm:ss]	Memory [GB]	Δt [hh:mm:ss]	Memory [GB]	Δt [hh:mm:ss]	Memory [GB]
bddbdb	0:30:00	5.7	DNF	DNF	DNF	DNF
SQLite	6:20:00	40.2	DNF	DNF	DNF	DNF
μZ	DNF	DNF	DNF	DNF	DNF	DNF
SOUFFLÉ	0:00:35	8.5	6:44:08	206.4	14:45:01	75.3

hand-crafted version of this analysis is reported to run under a minute [6]. For the CI analysis, bddbldb performs the analysis in a reasonable amount of time. However SOUFFLÉ outperforms bddbldb in terms of run-time by more than $50\times$ consuming a comparable amount of memory. In the case of the CS and Security analyses, SOUFFLÉ is the only tool capable of performing the analyses within the 18 h time limit imposed by the computation resources available to us for our evaluations. The Z3 based versions did not manage to finish any of our evaluated analyses in time.

4 Conclusion and Current Developments

We have presented SOUFFLÉ, a Datalog-based analysis tool that instead of evaluating Datalog, performs several specialization and optimization steps to produce a compiled, binary analyzer that can handle very large code bases. SOUFFLÉ is publicly available¹ and is actively developed by both Oracle and several universities. SOUFFLÉ supports a range of Datalog language extensions to aid in the specification of program analyses and resulting analyzers may be directly included into host applications as a header-only library.

Acknowledgement. We would like to thank Cristina Cifuentes, Paddy Krishnan, and all our other colleagues from Oracle Labs, Brisbane. We would also like to thank Byron Cook, Yannis Smaragdakis, and our anonymous reviewers.

References

1. Abiteboul, S., Hull, R., Vianu, V.: *Foundations of Databases*. Addison-Wesley, Boston (1995)
2. Allen, N., Krishnan, P., Scholz, B.: Combining type-analysis with points-to analysis for analyzing java library source-code. In: Møller, A., Naik, M. (eds.) *Proceedings of the 4th ACM SIGPLAN International Workshop on State of the Art in Program Analysis, SOAP@PLDI 2015, Portland, OR, USA, 15–17 June 2015*, pp. 13–18. ACM (2015)
3. Allen, N., Scholz, B., Krishnan, P.: Staged points-to analysis for large code bases. In: Franke, B. (ed.) *CC 2015. LNCS*, vol. 9031, pp. 131–150. Springer, Heidelberg (2015)
4. Alpuente, M., Feliú, M.A., Joubert, C., Villanueva, A.: Datalog-based program analysis with BES and RWL. In: de Moor, O., Gottlob, G., Furche, T., Sellers, A. (eds.) *Datalog 2010. LNCS*, vol. 6702, pp. 1–20. Springer, Heidelberg (2011)
5. Cifuentes, C., Gross, A., Keynes, N.: Understanding caller-sensitive method vulnerabilities: a class of access control vulnerabilities in the java platform. In: *Proceedings of the 4th ACM SIGPLAN International Workshop on State of the Art in Program Analysis. SOAP 2015, NY, USA*, pp. 7–12. ACM, New York (2015)
6. Dietrich, J., Hollingum, N., Scholz, B.: Giga-scale exhaustive points-to analysis for java in under a minute. In: *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications. OOPSLA 2015, NY, USA*. pp. 535–551. ACM, New York (2015)

¹ <http://souffle-lang.github.io>.

7. Dilworth, R.: A decomposition theorem for partially ordered sets. *Ann. Math.* **2**(51), 161–166 (1950)
8. Futamura, Y.: Partial evaluation of computation process - an approach to a compiler-compiler. *High. Order Symbolic Comput.* **12**(4), 381–391 (1999)
9. Green, T.J., Huang, S.S., Loo, B.T., Zhou, W.: Datalog and recursive query processing. *Found. Trends Databases* **5**(2), 105–195 (2013)
10. Hoder, K., Bjørner, N., de Moura, L.: μZ - an efficient engine for fixed points with constraints. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 457–462. Springer, Heidelberg (2011)
11. LogicBlox Inc.: Declarative cloud platform for applications that combine transactions & analytics. <http://www.logicblox.com>
12. Naik, M., Aiken, A., Whaley, J.: Effective static race detection for Java. *SIGPLAN Not.* **41**(6), 308–319 (2006)
13. Scholz, B., Jordan, H., Subotic, P., Westmann, T.: On fast large-scale program analysis in datalog. In: Zaks, A., Hermenegildo, M.V. (eds.) *Proceedings of the 25th International Conference on Compiler Construction, CC 2016, Barcelona, Spain, 12–18 March 2016*, pp. 196–206. ACM (2016)
14. Scholz, B., Vorobyov, K., Krishnan, P., Westmann, T.: A datalog source-to-source translator for static program analysis: an experience report. In: *24th Australasian Software Engineering Conference, ASWEC 2015, Adelaide, SA, Australia, 28 September – 1 October, 2015*, pp. 28–37. IEEE Computer Society (2015)
15. Smaragdakis, Y., Bravenboer, M., Kastrinis, G.: Doop: A framework for java pointer analysis. <http://doop.program-analysis.org/>
16. Smaragdakis, Y., Kastrinis, G., Balatsouras, G.: Introspective analysis: context-sensitivity, across the board. In: *PLDI, NY, USA, pp. 485–495*. ACM, New York (2014)
17. Veldhuizen, T.L.: C++ templates as partial evaluation. In: Danvy, O. (ed.) *PEPM*, pp. 13–18. University of Aarhus (1999). <http://dblp.uni-trier.de/db/conf/pepm/pepm1999.html#Veldhuizen99>
18. Whaley, J., Avots, D., Carbin, M., Lam, M.S.: Using datalog with binary decision diagrams for program analysis. In: Yi, K. (ed.) *APLAS 2005*. LNCS, vol. 3780, pp. 97–118. Springer, Heidelberg (2005)