

## Accurate Profiling in the Presence of Dynamic Compilation

Yudi Zheng<sup>1</sup>, Lubomír Bulej<sup>2</sup>, Walter Binder<sup>3</sup>

**Abstract:** Many programming languages are implemented on top of a managed runtime system, such as the Java Virtual Machine (JVM) or the .NET CLR, featuring an optimizing dynamic (just-in-time) compiler. Programs written in those languages are first interpreted (or compiled by a baseline compiler), whereas frequently executed methods are later compiled by the optimizing dynamic compiler.

Common feedback-directed optimizations [AHR02] performed by state-of-the-art dynamic compilers, such as the optimizing compiler in the Jikes RVM [Ar00] or Graal [Op], include method inlining and stack allocation of objects based on (partial) escape analysis [Ch99, SWM14], amongst others. Such optimizations result in compiled machine code that does not perform certain operations present at the bytecode level. In the case of inlining, method invocations are removed. In the case of stack allocation, heap allocations are removed and pressure on the garbage collector is reduced.

Many profiling tools are implemented using bytecode instrumentation techniques, inserting profiling code into programs at the bytecode level. However, because dynamic compilation is transparent to the instrumented program, a profiler based on bytecode instrumentation techniques is not aware of the optimizations performed by the dynamic compiler. Prevailing profilers based on bytecode instrumentation suffer from two serious limitations: (1) *over-profiling* of code that is optimized (and in the extreme case completely removed) by the dynamic compiler, and (2) *perturbation* of the compiler optimizations due to the inserted instrumentation code.

We present a novel technique to make profilers implemented with bytecode instrumentation techniques aware of the optimization decisions of the dynamic compiler, and to make the dynamic compiler aware of inserted profiling code. Our technique enables profilers which collect dynamic metrics that (1) correspond to an execution of the base program without profiling (w.r.t. the applied compiler optimizations), and (2) properly reflect the impact of dynamic compiler optimizations.

We implement our approach in a state-of-the-art Java virtual machine and demonstrate its significance with concrete profilers. We quantify the impact of escape analysis on allocation profiling, object lifetime analysis, and the impact of method inlining on callsite profiling. We illustrate how our approach enables new kinds of profilers, such as a profiler for non-inlined callsites, and a testing framework for locating performance bugs in dynamic compiler implementations.

This work was originally presented at OOPSLA'15 [ZBB15], where it received a *Distinguished Paper Award* as well as an endorsement from the Artifact Evaluation Committee for having submitted an easy-to-use, well-documented, consistent, and complete artifact<sup>4</sup>. In the meantime, the work has been integrated into the Graal project.

---

<sup>1</sup> Università della Svizzera italiana (USI), Faculty of Informatics, Switzerland, yudi.zheng@usi.ch

<sup>2</sup> Università della Svizzera italiana (USI), Faculty of Informatics, Switzerland, lubomir.bulej@usi.ch; Charles University, Faculty of Mathematics and Physics, Czech Republic

<sup>3</sup> Università della Svizzera italiana (USI), Faculty of Informatics, Switzerland, walter.binder@usi.ch

<sup>4</sup> <http://dag.inf.usi.ch/software/prof.acc/>

## Acknowledgments

The research presented here has been supported by Oracle (ERO project 1332), by the Swiss National Science Foundation (project 200021\_141002), by the European Commission (contract ACP2-GA-2013-605442), and by the Charles University institutional funding (SVV). We especially thank Thomas Würthinger and Lukas Stadler for their support with Graal.

## References

- [AHR02] Arnold, Matthew; Hind, Michael; Ryder, Barbara G.: Online Feedback-directed Optimization of Java. In: Proc. 17th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications. OOPSLA '02. ACM, pp. 111–129, 2002.
- [Ar00] Arnold, Matthew; Fink, Stephen; Grove, David; Hind, Michael; Sweeney, Peter F.: Adaptive Optimization in the Jalapeño JVM. In: Proc. 15th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications. OOPSLA '00. ACM, pp. 47–65, 2000.
- [Ch99] Choi, Jong-Deok; Gupta, Manish; Serrano, Mauricio; Sreedhar, Vugranam C.; Midkiff, Sam: Escape Analysis for Java. In: Proc. 14th ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications. OOPSLA '99. ACM, pp. 1–19, 1999.
- [Op] OpenJDK: , The Graal Compiler Project. <http://openjdk.java.net/projects/graal/>.
- [SWM14] Stadler, Lukas; Würthinger, Thomas; Mössenböck, Hanspeter: Partial Escape Analysis and Scalar Replacement for Java. In: Proc. IEEE/ACM International Symposium on Code Generation and Optimization. CGO '14. ACM, pp. 165:165–165:174, 2014.
- [ZBB15] Zheng, Yudi; Bulej, Lubomír; Binder, Walter: Accurate Profiling in the Presence of Dynamic Compilation. In: Proc. 30th ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications. OOPSLA '15. ACM, pp. 433–450, 2015.