Institut für Systemsoftware

O.Univ.Prof. Dr. Hanspeter Mössenböck

JOHANNES KEPLER
UNIVERSITÄT LINZ
Netzwerk für Forschung, Lehre und Praxis

Improving Truffle Interpreter Speed and Startup

Master thesis for …
Matr.Nr.: …
Email: …@...

The Truffle framework [1, 2] enables simple and efficient language implementation based on abstract-syntax-tree (AST) interpreters. Together with the use of the Graal just-in-time compiler, language implementations can reach the speed of state-of-the-art virtual machines. However, interpreter speed and startup performance have not yet been tackled.

One major issue of Truffle interpreters is their reliance on fine-grained nodes that can specialize during execution. While this results in great opportunities for optimizing peak-performance, it can have a significant negative impact on interpreter performance. Other design aspects such as the reliance on a generic mechanism to represent even small executable elements as methods come with a performance cost.

In this thesis, we want to explore whether ideas such as SuperInstructions [3] to coarsen the granularity of Truffle nodes based on usage patterns can be used to improve interpreter speed. Furthermore, we want to explore whether Truffle can be extended with new language-independent mechanisms such as AST-level inlining that works on the node level and does not stop at method boundaries.

The scope of this thesis is as follows:

- Develop a mechanism to create combinations of Truffle nodes based on the dynamic usage patterns to reduce overhead of guards and virtual dispatches in the interpreter.
- Explore language-independent inlining and related techniques to reduce the interpretation overhead without reducing peak performance.

The work's progress should be discussed with the supervisor at least every 2 weeks. Please note the guidelines of the Institute for System Software when preparing the written thesis.

Supervisor: Dr. Stefan Marr

[1] http://ssw.uni-linz.ac.at/Research/Projects/JVM/Truffle.html

[2] Würthinger, T., Wöß, A., Stadler, L., Duboscq, G., Simon D. and Wimmer, C. "Self-Optimizing AST Interpreters." In Proc. of the 8th Dynamic Languages Symposium, 2012.

[3] Casey, K.; Ertl, M. A. & Gregg, D. (2007), 'Optimizing Indirect Branch Prediction Accuracy in Virtual Machine Interpreters', ACM Trans. Program. Lang. Syst. 29 (6), 37. https://www.cs.tcd.ie/David.Gregg/papers/toplas05.pdf