Optimizing High Abstraction-Level Interpreters*

Stefan Brunthaler
Institut für Computersprachen
Technische Universität Wien
stefan@brunthaler.net

April 14, 2009

Abstract

Many users and companies alike feel uncomfortable with execution performance of interpreters, often also dismissing their use for specific projects. Specifically virtual machines whose abstraction level is higher than that of the native machine they run on, have performance issues. Several common existing optimization techniques fail to deliver their full potential on such machines. This paper presents an explanation for this situation and provides hints on possible alternative optimization techniques, which could very well provide substantially higher speedups.

Keywords: Interpreters, Virtual-Machine Abstraction, Optimization Techniques.

---

*This paper extends previous work, which has recently been presented at the *4th International Workshop on Bytecode Semantics, Verification, Analysis and Transformation* (*BYTECODE'09*), York, UK, 29th March 2009, part of ETAPS 2009.