

UNIVERSITY OF CALIFORNIA
Santa Barbara

Analysis, Detection, and Exploitation of Phase Behavior in Java Programs

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

by

Priya Nagpurkar

Committee in Charge:

Professor Chandra Krintz, Chair

Professor Timothy Sherwood

Professor Tobias Hollerer

September 2007

UMI Number: 3283656

PREVIEW

UMI[®]

UMI Microform 3283656

Copyright 2008 by ProQuest Information and Learning Company.
All rights reserved. This microform edition is protected against
unauthorized copying under Title 17, United States Code.

ProQuest Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

The Dissertation of
Priya Nagpurkar is approved:

Professor Timothy Sherwood

Professor Tobias Hollerer

Professor Chandra Krintz, Committee Chairperson

August 2007

Analysis, Detection, and Exploitation of Phase Behavior in Java Programs

Copyright © 2007

by

Priya Nagpurkar

PREVIEW

Dedication and Gratitude

I dedicate this dissertation to my parents, Ashok and Rekha Nagpurkar, who never cease to amaze me with the perfect balance between independence and guidance that they have always struck in influencing my life. Their faith, encouragement and support, further reinforced by that of my brother, Ravindra, has played an important role in my decision to pursue, and in successfully completing this work.

My advisor, Chandra Krintz, played an equally important role, by recognizing and bringing out my potential for research. Her energy and enthusiasm for both research, and teaching were, and will continue to be, a constant inspiration. She has been a model *guru*, by being a good friend, philosopher, and guide. Many thanks also to Tim Sherwood, and Tobias Hoellerer for their guidance as members of my dissertation committee, and to our collaborators from the I.B.M. T.J. Watson Research Center for their valuable advice. Michael Hind, Peter Sweeney, Trey Cain, Mauricio Serrano, and Jong-Deok Choi were all excellent mentors.

I would like to express deep gratitude towards all my friends, old and new. Special thanks to my close friends, Rekha, Shilpa, Puja, Nicole, and Martina for being my extremely reliable support structure in times of need; to Lingli and Ye, for being great colleagues and neighbors; to Selim for great evenings in the climbing gym; and to Hussam for all the cups of tea; These and other members of the RACE lab made it a great place to work or to hang out. Visits to the CS office were always pleasant, thanks to Amanda, Greta, Julia, Beejay, and the rest of our very cheerful office staff. Finally, playing ultimate with the CS team was always something to look forward to – thanks to all of you on the ultimate team for your camaraderie!

Acknowledgements

The text of Chapter 3 is in part a reprint of the material as it appears in the proceedings of Elsevier Science of Computer Programming – Special Issue on Principles Practices and Programming in Java, Vol. 59. The dissertation author was the primary researcher and author and the co-author listed on this publication ([82]) directed and supervised the research which forms the basis for Chapter 3.

The text of Chapter 4 is in part a reprint of the material as it appears in the proceedings of the Fourth Annual International Symposium on Code Generation and Optimization (CGO). The dissertation author was the primary researcher and the co-authors listed on this publication ([80]) directed and supervised the research which forms the basis for Chapter 4.

The text of Chapter 6 is in part a reprint of the material as it appears in the proceedings of ACM Transactions on Architecture and Code Optimization (TACO), Vol. 3, Number 1. The dissertation author was the primary researcher and author with significant contribution from one of the co-authors, Hussam Mousa. The remaining co-authors listed on this publication ([84]) directed and supervised the research which forms the basis for Chapter 6.

The text of Chapter 7 is in part a reprint of the material as it appears in the proceedings of the Sixteenth International Conference on Parallel Architectures and Compilation Techniques (PACT). The dissertation author was the primary researcher and author and the co-authors listed on this publication ([79]) directed and supervised the research which forms the basis for Chapter 7.

Curriculum Vitæ

Priya Nagpurkar

Education

- 2007 **Doctor of Philosophy in Computer Science,**
University of California, Santa Barbara.
- 2007 **Master of Science in Computer Science,**
University of California, Santa Barbara.
- 2001 **Bachelor of Engineering in Computer Engineering,**
Pune University.

Professional Experience

- 2006 **Summer Intern,**
I.B.M. T.J. Watson Research Center.
- 2003 – 2007 **Graduate Research Assistant,**
University of California, Santa Barbara.
- 2001 – 2003 **Graduate Teaching Assistant,**
University of California, Santa Barbara.
- 2000 **Intern,**
VERITAS Software India Ltd. (now Symantec), Pune, India.

Professional Activities

- 2007 **Program Committee Member,** International Conference on Principles and Practices of Programming in Java
- 2006 **Program Committee Member,** UCSB Graduate Student Research Conference
- 2006 **Submissions Chair,** International Conference on Principles and Practices of Programming in Java
- 2006-2007 **Graduate Student Representative,** Colloquium Committee
- 2005-2006 **Graduate Student Representative,** Graduate Admissions Committee

Publications

Priya Nagpurkar, Harold W. Cain, Mauricio Serrano, Jong-Deok Choi and Chandra Krintz: “Call-chain Software Instruction Prefetching in J2EE Server Applications,” *In the Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT07)*

Lingli Zhang, Chandra Krintz, and Priya Nagpurkar: “Language and Virtual Machine Support for Efficient Fine-Grained Futures in Java,” *In the Proceedings of the International Conference on Parallel Architectures and Compilation Techniques (PACT07)*

Lingli Zhang, Chandra Krintz, and Priya Nagpurkar: “Supporting Exception Handling for Futures in Java,” *In the Proceedings of the International Conference on the Principles and Practice on Programming in Java (PPPJ07)*

Priya Nagpurkar, Harold W. Cain, Mauricio Serrano, Jong-Deok Choi and Chandra Krintz: “A Study of Instruction Cache Performance and the Potential for Instruction Prefetching in J2EE Server Applications,” *Tenth Workshop on Computer Architecture Evaluation Using Commercial Workloads (CAECW-10)*

Priya Nagpurkar, Michael Hind, Chandra Krintz, Peter Sweeney, and V.T. Rajan: “On-line Phase Detection Algorithms,” *In the Proceedings of the International Symposium on Code Generation and Optimization (CGO06)*

Priya Nagpurkar, Chandra Krintz, and Timothy Sherwood: “Phase-aware Remote Profiling,” *In the Proceedings of the International Symposium on Code Generation and Optimization (CGO05)*

Priya Nagpurkar and Chandra Krintz: “Visualization and Analysis of Phased Behavior in Java Programs,” *In the Proceedings of the International Conference on the Principles and Practice of Programming in Java (PPPJ04)*

Priya Nagpurkar, Hussam Mousa, Chandra Krintz, and Timothy Sherwood: “Efficient Remote Profiling for Resource-Constrained Devices,” *In the Proceedings of ACM Transactions on Architecture and Code Optimization (TACO), Vol. 3, Number 1, March, 2006, pages 1-32*

Priya Nagpurkar, and Chandra Krintz: “Phase-Based Visualization and Analysis of Java Programs,” *In the Proceedings of Elsevier Science of Computer Programming – Special Issue on Principles Practices and Programming in Java, Vol. 59, Number 1-2, January 2006, pages 64-81*

Selim Gurun, Priya Nagpurkar and Ben Zhao: “Energy Consumption and Conservation in Mobile Peer-to-peer Systems,” *The first International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking (MobiShare06)*

Field of Study: Computer Science

PREVIEW

Abstract

Analysis, Detection, and Exploitation of Phase Behavior in Java Programs

Priya Nagpurkar

The Java programming language offers developers many productivity enhancing features, including high-level abstractions, extensive libraries, architecture-independent execution, and type safety. These features are enabled by an intelligent execution environment that, incrementally and dynamically, compiles and executes compact representations of Java programs encoded for a virtual machine. While this necessarily adds overhead, the ability to compile (and recompile) code at runtime also enables the execution environment to perform dynamic, performance-enhancing optimizations based on the runtime behavior of the executing program. There are three primary steps in developing effective adaptive optimizations for these systems: (1) Development of a thorough analysis, understanding, and characterization of the performance of Java programs; (2) Extracting accurate data from programs efficiently at runtime; and (3) Guiding optimizations using feedback from the extracted performance data.

We address each of these steps in our research by focusing on techniques that capture and exploit the repeating patterns in program behavior (phases) within virtual execution environments, and in particular, those for Java programs. This dissertation can

be decomposed into two foci: phase analysis and detection tools and techniques and phase-aware techniques for efficient program analysis and optimization. We first study the time varying behavior of Java programs, show that Java programs do exhibit phase behavior, and present tools to extract and analyze this phase behavior. We then investigate the problem of accurate online phase detection for Java programs, within a Java virtual machine, the parameters that impact doing so effectively, and evaluate numerous online phase detectors. Finally we demonstrate the potential of phase-based optimizations by designing and evaluating two phase-based runtime techniques. The first technique is an accurate, low-overhead profiling scheme for resource-constrained devices that uses phases to drive when to sample the execution of a program. The second technique is a software instruction prefetching mechanism that uses method-level phase behavior to identify, predict, and prefetch methods that incur a large number of instruction cache misses for emerging Java workloads like database- and application servers. These two techniques span two extremes of execution environments used for Java applications: software for resource-constrained devices at the low end and application servers at the high-end.

Contents

Acknowledgements	v
Curriculum Vitæ	vi
Abstract	ix
List of Figures	xiv
List of Tables	xvi
1 Introduction	1
2 Background	6
2.1 Phase Characterization, Detection, and Prediction Techniques	7
2.2 Applications of Phase Analysis	16
2.3 Dynamic Compilation and Adaptive Optimization in Java	20
3 Phase Behavior in Java Programs	23
3.1 Phase Analysis Framework	25
3.1.1 Data Generation	26
3.1.2 Data Processing	29
3.2 Phase Analysis Toolkit	31
3.2.1 Phase Visualizer	31
3.2.2 Phase Finder	33
3.2.3 Phase Analyzer and Code Extractor	36
3.3 Analysis	37
3.3.1 Visual Analysis	38
3.3.2 Efficient Identification of Optimization Opportunities	44

3.3.3	Cross-Input Analysis	49
3.3.4	Other Opportunities for Exploiting Phase Behavior	51
3.4	Summary	53
4	Phase Detection for Java Programs	55
4.1	Online Phase Detection Framework	56
4.1.1	Window Policy	61
4.1.2	Model Policy	63
4.1.3	Analyzer Policy	64
4.2	Evaluating Phase Detectors	65
4.2.1	Phase Detection Baseline	66
4.2.2	Accuracy Scoring Metric	69
4.3	Analysis	71
4.3.1	Methodology	73
4.3.2	Window Policy	77
4.3.3	Model Policy	81
4.3.4	Analyzer Policy	84
4.3.5	Additional Analysis	85
4.4	Summary	90
5	Phase-based Runtime Techniques	91
5.1	Phase-aware Profiling	92
5.2	Instruction Prefetching	94
6	Phase-aware Remote Profiling	97
6.1	Phase-aware Sampling: Deciding When to Sample	101
6.2	Profiling Support for Toggling Profile Collection	106
6.2.1	Dynamic Instruction Stream Editing (DISE)	108
6.2.2	Hybrid Profiling Support using DISE	110
6.3	Evaluation	118
6.3.1	Phase-aware Profiling for General-purpose Programs	119
6.3.2	Phase-aware Profiling for Embedded Devices	133
6.4	Extending Phase-aware Profiling to Multiple Users	139
6.5	Related Work	143
6.5.1	Efficient Profiling	143
6.5.2	Monitoring Program Behavior for Bug Isolation and Test Coverage	145
6.6	Summary	147

7	Phase-based Instruction Prefetching	149
7.1	Characterization of Instruction Cache Behavior	150
7.1.1	Methodology	151
7.1.2	Stall Cycles	152
7.1.3	Method-level Analysis	154
7.2	Method-level Phase Behavior	158
7.3	Call-chain Instruction Prefetching	161
7.3.1	Design and Implementation	161
7.3.2	Experimental Methodology	165
7.3.3	Evaluation	167
7.3.4	Discussion: Potential Improvements	174
7.4	Related Work	175
7.5	Summary	178
8	Conclusion	179
8.1	Dissertation Summary	180
8.2	Impact and Future Directions	187
	Bibliography	191

List of Figures

1.1	Programming Language Usage Trends.	2
3.1	JVM phase analysis framework and toolkit	25
3.2	Architecture of the data generation framework	27
3.3	Phase Visualizer	32
3.4	Similarity graph for Mtrt input size 10.	39
3.5	Phases for Mtrt with similarity threshold 0.8	40
3.6	Similarity graphs for the SpecJVM benchmarks with input size 100	42
3.7	Similarity graphs for the SpecJVM benchmarks with input size 10	43
3.8	Code extracted using the phase framework and toolkit	47
3.9	Hand-optimized basic block exposed via phase analysis	48
3.10	Analysis of cross-input similarity	50
4.1	Illustrated view of the phase detection framework	57
4.2	Basic operation of the phase detection framework	58
4.3	Online phase detection framework	62
4.4	Evaluation of skip factor and Fixed versus Adaptive windowing	80
4.5	Unweighted vs. Weighted similarity models	82
4.6	Constant vs. Adaptive window policy	84
4.7	Slide vs. Move resizing	88
4.8	Accurate detection of phase boundaries	89
6.1	Run-time power usage	98
6.2	System overview	100
6.3	Overview of the phase-aware profiling scheme	102
6.4	The Hybrid Profiling Support (HPS) system	107
6.5	HPS extensions to DISE	112
6.6	HPS pattern and replacement specification grammar	113

6.7	Pattern and replacement productions for different profile types	118
6.8	DISE vs. HPS for performance sampling	122
6.9	Evaluation of representative selection policies	125
6.10	Average error in code region profiling	127
6.11	Efficacy across profile types	130
6.12	Evaluation of phase-aware sampling using the StrongARM environment and benchmarks	136
6.13	Distributed profiling across multiple executions	140
7.1	Commit Stall Cycle Categorization	152
7.2	Icache misses per 100 committed instructions	153
7.3	Per-method Contribution to Total icache Misses (cumulative distribution)	154
7.4	Method-level phases in WebSphere (running specjAppServer2001)	157
7.5	Correlation between method-level phases and phases in icache misses.	159
7.6	Overview of phase-based prefetching in a JVM.	160
7.7	Example call chain	162
7.8	Trace-based Analysis Methodology	166
7.9	Prefetch accuracy	170
7.10	Effect of miss distance on coverage and interference	172

PREVIEW

List of Tables

3.1	Description of the benchmarks used.	38
4.1	Benchmark Characteristics	72
4.2	Window size comparison	74
6.1	Select benchmark statistics relevant to the profiles collected	120
6.2	Sampling overhead at 5% error	132
6.3	StrongARM methodology	134
7.1	Prefetch Target Characteristics.	156
7.2	Coverage achieved	171
7.3	Improvement in IPC	173

Chapter 1

Introduction

The greatest happiness for the thinking person is to have explored the explorable and to venerate in equanimity that which cannot be explored.

Johann Wolfgang von Goethe (1749-1832)

The Java programming language, and similarly C# and the Microsoft .Net languages, offer many benefits to programmers such as portability, programmer productivity through high-level abstractions and extensive libraries, type and memory safety, and dynamic loading. These features make it easier, not only to develop software, but also to debug and maintain it. As a result, these languages are very popular with software developers. Java, in particular, has seen tremendous growth since its inception, a little over a decade ago, and the trend is predicted to continue [27]. It is estimated that Java today drives a \$100 billion a year software industry, and is deployed on a wide variety of devices, including millions of desktops, billions of embedded devices (from smart phones to car navigation systems), and enterprise servers [101, 1].

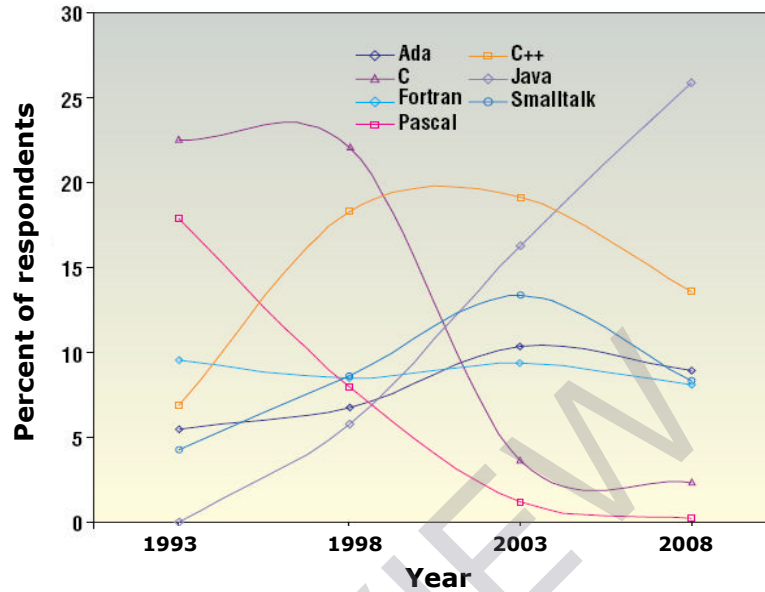


Figure 1.1: Programming Language Usage Trends. This figure from [27] uses historic data, gathered from real users, as well as prediction based on this data to show the long term trend in programming language popularity.

To enable these features, especially portability (the write-once, run-anywhere model), these programs are compiled into an architecture-independent intermediate format and executed within a virtual execution environment on the target host. The execution environment, a Java virtual machine or .Net runtime, implements a compilation system that converts the intermediate code to the native format of the underlying machine. This dynamic compilation necessarily introduces runtime overhead, but at the same time also exposes opportunities for adaptation – optimizations that we can customize according to the behavior of the executing program. State of the art Java Virtual Machines (JVMs) [71, 28, 102, 5, 53] employ adaptive optimization techniques based on

information gathered while the program is executing (feedback-directed optimization). Monitoring, analyzing, and predicting runtime program behavior are vital to feedback-directed optimization.

Recent research has shown, for non-Java programs, that program behavior varies over time, and exhibits repeating patterns [93, 41], and has focused on automatically characterizing this behavior [94, 95, 66, 35]. Phase analysis of programs is one such characterization, which isolates distinct behaviors in a program's execution by grouping periods of execution that are similar together in a *phase*. A program's execution can then be seen as a series of phases that might repeat themselves several times. Researchers have used phase behavior to improve program performance via hardware and software optimization [38, 96, 92, 72] and to reduce simulation time [94, 95] and generate cycle-close traces [88]. The ability to detect and predict phases at runtime has the potential of uncovering new opportunities in performing proactive adaptive optimizations for Java programs. With the ultimate aim of enabling better performance for Java programs, the thesis question that we explore is this work is:

How can we efficiently detect, track, predict, and exploit repeating patterns in program behavior (phases) in a Java virtual machine to facilitate runtime analysis and feedback-directed optimization in Java programs?

To answer this question, we focus our efforts on phase characterization and detection in Java programs, and exploiting repetitions that phases manifest, using optimiza-

tion for a range of devices for which Java Virtual Machines are available. In particular, we

- investigate and develop offline mechanisms and tools that enable the characterization, visualization, and manipulation of phase behavior in Java programs,
- develop a modular, pluggable framework for implementing and investigating on-line phase detection algorithms within a Java Virtual Machine,
- devise and investigate a phase-aware approach to collecting accurate online execution profiles
- devise and investigate a prefetching scheme (a dynamic optimization) that exploits repeating patterns in Java server execution.

We take an empirical and implementation-oriented approach to developing phase detection and exploitation techniques in this thesis. We implement and empirically evaluate our techniques using a wide range of real programs and open-source Java virtual machine technologies. From this effort, we have produced a set of tools that significantly facilitate analysis of repeating patterns in the behavior of Java programs, and we have designed and evaluated novel techniques for exploiting phases to improve program profiling and execution performance.

The dissertation is organized as follows. We begin with a discussion of the relevant background in Chapter 2. This chapter includes techniques that we use in our work, as

well as the state of the art for phase analysis and its uses. We present our framework and toolkit for understanding and analyzing phase behavior in Java programs in Chapter 3, followed by our framework for the design and analysis of online phase detection algorithms within a JVM, in Chapter 4. Chapters 5, 6, and 7 focus on the use of phase behavior to enable two phase-based runtime techniques; Chapter 5 introduces these techniques and provides an overview, while the following Chapters provide details of each technique. We end with a summary of the dissertation and a discussion of future directions in Chapter 8.

PREVIEW

Chapter 2

Background

The focus of this dissertation is understanding, analyzing, and exploiting the repeating patterns, or phases, in the time varying behavior of Java programs. Of particular interest to us, is the possibility of incorporating phase-awareness in virtual execution environments, like Java virtual machines, to drive adaptive, feedback-directed optimization of dynamically compiled programs.

Much research has already gone into the characterization, detection, prediction, and exploitation of phase behavior, especially in the area of computer architecture. In this Chapter, we present an overview of extant work on phase behavior, and also briefly describe the process of dynamic compilation for Java.

2.1 Phase Characterization, Detection, and Prediction Techniques

Program behavior has commonly been abstracted in the form of profiles gathered over the program's execution. Recently, there has been a lot of interest in studying program behavior during different parts of execution. Many researchers have observed, through detailed simulations and temporal profiles, that programs exhibit widely varying behavior during different parts of execution [93, 94, 41]. Program behavior, however, is not entirely random and often shows significant structure. In [94] and [41], the authors periodically gathered various hardware metrics, like IPC, cache misses, branch misprediction rate with the aim of studying low-level program behavior over time and finding any possible correlation between the metrics. Their findings indicate that, not only does program behavior change, but it also has periods of stable execution interspersed with transitions. During periods of stable execution, the architectural metrics measured are relatively stable. What is more interesting is the fact that the metrics transition in unison, though the nature of the transition might be different (that is the instruction cache miss rate might go up, whereas the IPC might go down). Recognizing the importance of automatically characterizing this behavior in order to exploit it for various purposes (like reducing simulation time, aiding prediction and proactive optimization), various techniques were developed at different levels in the system stack –