ARCHITECTURAL SUPPORT FOR ENHANCED PERFORMANCE OF OS INTENSIVE APPLICATIONS

PRATHMESH KALLURKAR

DEPARTMENT OF COMPUTER SCIENCE & ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY DELHI
JANUARY 2018

ARCHITECTURAL SUPPORT FOR ENHANCED PERFORMANCE OF OS INTENSIVE APPLICATIONS

by

PRATHMESH KALLURKAR

Department of Computer Science and Engineering

Submitted

in fulfillment of the requirements of the degree of Doctor of Philosophy

to the

Indian Institute of Technology Delhi

JANUARY 2018

Certificate

This is to certify that the thesis titled **ARCHITECTURAL SUPPORT FOR ENHANCED PERFORMANCE OF OS INTENSIVE APPLICATIONS** being submitted by **Mr. PRATHMESH KALLURKAR** for the award of **Doctor of Philosophy** in Computer Science and Engineering is a record of bona fide work carried out by him under my guidance and supervision at the Department of Computer Science and Engineering, Indian Institute of Technology Delhi. The work presented in this thesis has not been submitted elsewhere, either in part or full, for the award of any other degree or diploma.

Smruti Ranjan Sarangi
Associate Professor
Department of Computer Science and Engineering
Indian Institute of Technology Delhi
New Delhi- 110016

Acknowledgements

Firstly, I would like to express my sincere gratitude towards my supervisor Dr. Smruti Sarangi for his continuous support and guidance during the entire duration of my PhD. He encouraged me to take a plunge in this research area around seven years back, and I must admit that this was one of the best decision I have made till now. His valuable suggestions helped me improve the quality of research. He also spent a significant amount of time in improving my presentation and writing skills. Thank you Sir.

I thank Dr. Anshul Kumar, Dr. Preeti Ranjan Panda, and Dr. G.S. Visweswaran for providing me valuable feedback and insights on my research work. I would also like to thank Dr. Sorav Bansal for teaching me Operating systems and Virtualization.

I thank all my friends at IIT Delhi. Your constant support and encouragement ensured that I could navigate the highs and lows of life with relative ease. The experiences that we shared were thoroughly enjoyable; I will cherish them forever.

I also thank my parents and my sister for all the constant support and guidance during my PhD. Thank you Aai, Baba, and Prajkta.

Prathmesh Kallurkar

Abstract

The OS has a large impact on the performance of important workloads such as web servers, database servers, and file servers. We observe that such workloads spend a considerable portion (20-95%) of their time executing different OS tasks such as interrupt handlers, system call handlers, and the OS scheduler; hence, we call them as OS-intensive workloads. We start by proposing a full-system simulation framework to do an in-depth study of such full-system workloads. Using the developed simulation framework, we show that the instruction footprint of such workloads is typically much larger than the size of instruction caches (16-32 KB). We also bring out interesting insights regarding the execution of such workloads. Based on the detailed insights regarding such workloads, we propose three novel techniques to improve the performance of OS-intensive workloads: (1) an interference-aware instruction cache, (2) an instruction prefetcher that is optimized for OS-intensive applications, and (3) a task scheduler that is optimized for OS-intensive applications.

सार

वेब सर्वर, डेटाबेस सर्वर, और फ़ाइल सर्वर जैसे महत्वपूर्ण वर्कलोड के प्रदर्शन पर ओएस का बड़ा प्रभाव पड़ता है। हम देखते हैं कि ऐसे वर्कलोड्स अलग-अलग ओएस कार्यों जैसे इंटरप्ट हैंडलर, सिस्टम कॉल हैंडलर और ओएस शेड्यूलर को कार्यान्वयन करने में अपने समय का एक बड़ा हिस्सा (बीस से पचानवे प्रतिशत) खर्च करते हैं; इसलिए, हम उन्हें ओएस-गहन वर्कलोड कहते हैं। हम इस तरह के पूर्ण-सिस्टम वर्कलोड के गहन अध्ययन करने के लिए एक पूर्ण-प्रणाली सिमुलेशन फ्रेमवर्क का प्रस्ताव देकर शुरू करते हैं। विकसित सिमुलेशन फ्रेमवर्क का उपयोग करके, हम दिखाते हैं कि ऐसे वर्कलोड का निर्देश पदचिह्न आमतौर पर निर्देश कैश (सोलह से बत्तीस किलो बाइट) के आकार से काफी बड़ा होता है। हम ऐसे वर्कलोड के कार्यान्वयन के संबंध में दिलचस्प अंतर्दृष्टि भी लाते हैं। ऐसे वर्कलोड के बारे में विस्तृत अंतर्दृष्टि पर आधारित, हम ओएस-गहन वर्कलोड के प्रदर्शन को बेहतर बनाने के लिए तीन उपन्यास तकनीकों का प्रस्ताव देते हैं (१) एक हस्तक्षेप-जागरूक निर्देश कैश, (२) एक निर्देश प्रीफेचर जिसे ओएस-गहन अनुप्रयोगों के लिए अनुकूलित किया गया है, और (३) एक कार्य शेड्यूलर जिसे ओएस-गहन अनुप्रयोगों के लिए अनुकूलित किया गया है।

Contents

List of Figures

List of Tables