Journal of Instruction-Level Parallelism 10 (2008) 1-2 Submitted 10/07; published 6/08

Best Papers from ACM's Computing Frontiers Conference

Sally A. McKee SAM@CSL.CORNELL.EDU
Computer Systems Laboratory
Department of Electrical and Computer Engineering
324 Rhodes Hall
Cornell University
Ithaca, NY 14853 USA

1. Introduction

This Special Section consists of selected papers from ACM SIGMICRO's Computing Frontiers (CF) Conference, held each May in Ischia, Italy. Most papers appeared in CF 2006, but one related paper (on "pipeline spectroscopy") appeared in CF 2007.

2. Contents

The first paper, "On the Nature of Cache Miss Behavior: Is It sqrt(2)?" imparts insight into a rule of thumb that many of us have long used. The second, related paper, "Analyzing the Cost of a Cache Miss Using Pipeline Spectroscopy", likewise attempts to yield insight, where most cache studies focus on data, rather than information. Both come from IBM T.J. Watson Research Center.

The third and fourth papers come from Washington University in St. Louis. The third, "Dynamic Thread Assignment on Heterogeneous Multiprocessor Architectures", discusses productive use of heterogeneous CMPs, and the fourth "Exploiting Locality to Ameliorate Packet Queue Contention and Serialization" introduces a queuing cache for network processing: a hardware cache and a closely coupled queuing engine implementing queuing operations moves the bottleneck from external memory onto packet processors with lower latencies.

The fifth paper, "Dynamic Parallelization and Mapping of Binary Executables on Hierarchical Platforms", presents a software-based solution to generate parallel applications for any parallel hardware substrate, obviating the need to hand-parallelize each application for each new hardware platform.

The sixth paper, "Refactoring Intermediately Executed Code to Reduce Cache Capacity Misses", presents another software technique that modifies executables, this time to improve memory hierarchy performance. Likewise, the seventh paper, "Simple Penalty-Sensitive Cache Replacement Policies", seeks to improve memory hierarchy performance, but through more intelligent management policies.

The last paper addresses a very different kind of optimization in an area significant for much of the software community: "Instruction Folding in a Hardware Translation-Based Java Virtual Machine". This solution is elegant and creative, performing well at low hardware complexity.

Acknowledgments

I thank the JILP Editor in Chief, Eric Rotenberg, and Associate Editor in Chief, Dan Connors, for allowing me to present this special section, and for their patience in my producing it. Other responsibilities delayed delivery of the included material in the required format.

Finally, I thank the many contributors to and participants in Computing Frontiers. It is truly a unique meeting, fostering the exchange of new ideas, the debate of existing ideas, the enjoyment of new and old friends alike, and (last but not least) the enjoyment of a great deal of culture, both intellectual and personal. I thank the many dedicated members of the CF Steering Committee, and particularly Monica Alderighi, Sergio D'Angelo, and Gerry Johnson for first involving me with this wonderful group of people.

This set of papers is dedicated to the memory of Prof. Stamatis Vassiliadis from TU-Delft, without whose contributions and enthusiasm Computing Frontiers would not have evolved as it has. He is sorely missed.