

2008 International Conference on Parallel Architectures and Compilation Techniques (PACT 2008)

**Toronto, Ontario, Canada
25 – 29 October 2008**



**IEEE Catalog Number: CFP08073-POD
ISBN: 978-1-5090-3021-7**

**Copyright © 2008, The Association for Computing Machinery (ACM)
All Rights Reserved**

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP08073-POD
ISBN (Print-On-Demand):	978-1-5090-3021-7
ISBN (Online):	978-1-60558-282-5
ISSN:	1089-795X

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

Table of Contents

Organization	ix
---------------------------	----

Keynote

• GPU Evolution: Will Graphics Morph into Compute?	1
Norm Rubin (<i>Advanced Micro Devices</i>)	

Session 1: Compilation

• Outer-Loop Vectorization — Revisited for Short SIMD Architectures	2
Dorit Nuzman, Ayal Zaks (<i>IBM Haifa Research Laboratory</i>)	
• Redundancy Elimination Revisited	12
Keith Cooper, Jason Eckhardt, Ken Kennedy (<i>Rice University</i>)	
• Exploiting Loop-Dependent Stream Reuse for Stream Processors	22
Xuejun Yang, Ying Zhang (<i>National University of Defence Technology, China</i>), Jingling Xue (<i>The University of New South Wales</i>), Ian Rogers (<i>The University of Manchester</i>), Gen Li, Guibin Wang (<i>National University of Defence Technology, China</i>)	
• Feature Selection and Policy Optimization for Distributed Instruction Placement Using Reinforcement Learning	32
Katherine E. Coons, Behnam Robatmili, Matthew E. Taylor, Bertrand A. Maher, Doug Burger, Kathryn S. McKinley (<i>University of Texas at Austin</i>)	

Session 2: CMP Architecture Design

• Core Cannibalization Architecture: Improving Lifetime Chip Performance for Multicore Processors in the Presence of Hard Faults	43
Bogdan F. Romanescu, Daniel J. Sorin (<i>Duke University</i>)	
• Pangaea: A Tightly-Coupled IA32 Heterogeneous Chip Multiprocessor	52
Henry Wong (<i>University of British Columbia</i>), Anne Bracy, Ethan Schuchman (<i>Intel Corporation</i>), Tor M. Aamodt (<i>University of British Columbia</i>), Jamison D. Collins, Perry H. Wang, Gautham China, Ankur Khandelwal Groen, Hong Jiang, Hong Wang (<i>Intel Corporation</i>)	
• Skewed Redundancy	62
Gordon B. Bell (<i>IBM Corporation</i>), Mikko H. Lipasti (<i>University of Wisconsin - Madison</i>)	

Session 3A: Analyzing Applications

• The PARSEC Benchmark Suite: Characterization and Architectural Implications	72
Christian Bienia (<i>Princeton University</i>), Sanjeev Kumar (<i>Intel Corporation</i>), Jaswinder Pal Singh, Kai Li (<i>Princeton University</i>)	
• Visualizing Potential Parallelism in Sequential Programs	82
Graham D. Price, John Giacomoni, Manish Vachharajani (<i>University of Colorado</i>)	
• Characterizing and Modeling the Behavior of Context Switch Misses	91
Fang Liu, Fei Guo, Yan Solihin (<i>North Carolina State University</i>), Seongbeom Kim (<i>VMWare Inc.</i>), Abdulaziz Eker (<i>The Scientific and Technical Research Council</i>)	

Session 3B: I/O Optimizations

• MCAMP: Communication Optimization on Massively Parallel Machines with Hierarchical Scratch-pad Memory	102
Hiroshige Hayashizaki, Yutaka Sugawara, Mary Inaba, Kei Hiraki (<i>The University of Tokyo</i>)	
• Profiler and Compiler Assisted Adaptive I/O Prefetching for Shared Storage Caches	112
Seung Woo Son, Sai Prashanth Muralidhara (<i>Pennsylvania State University</i>), Ozcan Ozturk (<i>Bilkent University</i>), Mahmut Kandemir (<i>Pennsylvania State University</i>), Ibrahim Kolcu (<i>University of Manchester</i>), Mustafa Karakoy (<i>Imperial College</i>)	
• Runtime Optimization of Vector Operations on Large Scale SMP Clusters	122
Costin Iancu, Steven Hofmeyr (<i>Lawrence Berkeley National Laboratory</i>)	

Keynote

- **(How) Can Programmers Conquer the Multicore Menace?**N/A
Saman Amarasinghe (*Massachusetts Institute of Technology*)

Session 4: Multicore Memory Hierarchy Design (Part 1)

- **Distributed Cooperative Caching** 134
Enric Herrero, José González, Ramon Canal (*Universitat Politècnica de Catalunya*)
- **Scalable and Reliable Communication for Hardware Transactional Memory** 144
Seth H. Pugsley, Manu Awasthi, Niti Madan, Naveen Muralimanohar,
Rajeev Balasubramonian (*University of Utah*)
- **Improving Support for Locality and Fine-Grain Sharing in Chip Multiprocessors** 155
Hemayet Hossain, Sandhya Dwarkadas, Michael C. Huang (*University of Rochester*)

Session 5: Reconfigurable Architecture Optimization

- **Edge-centric Modulo Scheduling for Coarse-Grained Reconfigurable Architectures** 166
Hyunchul Park, Kevin Fan, Scott A. Mahlke (*University of Michigan*),
Taewook Oh, Heeseok Kim, Hong-seok Kim (*Samsung Advanced Institute of Technology*)
- **Multi-Optimization Power Management for Chip Multiprocessors** 177
Ke Meng, Russ Joseph, Robert P. Dick (*Northwestern University*), Li Shang (*University of Colorado*)
- **Multitasking Workload Scheduling on Flexible-Core Chip Multiprocessors** 187
Divya P. Gulati (*University of Texas at Austin*), Changkyu Kim (*Intel Corporation*),
Simha Sethumadhavan (*Columbia University*), Stephen W. Keckler, Doug Burger (*University of Texas at Austin*)

Session 6: Multicore Memory Hierarchy Design (Part 2)

- **Leveraging On-Chip Networks for Data Cache Migration in Chip Multiprocessors** 197
Noel Eisle, Li-Shiuan Peh (*Princeton University*), Li Shang (*University of Colorado*)
- **Adaptive Insertion Policies for Managing Shared Caches** 208
Aamer Jaleel, William Hasenplaugh (*Intel Corporation*), Moinuddin Qureshi (*IBM T.J. Watson Research Center*),
Julien Sebot (*Intel Israel Design Center*), Simon Steely Jr., Joel Emer (*Intel Corporation*)
- **Analysis and Approximation of Optimal Co-Scheduling on Chip Multiprocessors** 220
Yunlian Jiang, Xipeng Shen (*College of William and Mary*),
Jie Chen (*Thomas Jefferson National Accelerator Facility*), Rahul Tripathi (*University of South Florida*)

Session 7: Multithreading Improvements

- **An Adaptive Resource Partitioning Algorithm for SMT Processors** 230
Huaping Wang, Israel Koren, C. Mani Krishna (*University of Massachusetts*)
- **Meeting Points: Using Thread Criticality to Adapt Multicore Hardware to Parallel Regions** 240
Qiong Cai, José González (*Intel Barcelona Research Center*), Ryan Rakvic (*United States Naval Academy*),
Grigorios Magklis, Pedro Chaparro, Antonio González (*Intel Barcelona Research Center*)

Session 8: Middleware and Runtime Systems

- **Prediction Models for Multi-dimensional Power-Performance Optimization on Many Cores** 250
Matthew Curtis-Maury, Ankur Shah, Filip Blagojevic, Dimitrios S. Nikolopoulos
(*Virginia Polytechnic Institute and State University*),
Bronis R. de Supinski, Martin Schulz (*Lawrence Livermore National Laboratory*)
- **Mars: A MapReduce Framework on Graphics Processors** 260
Bingsheng He, Wenbin Fang, Qiong Luo (*Hong Kong University of Science and Technology*),
Naga K. Govindaraju (*Microsoft Corporation*), Tuyong Wang (*Sina Corporation*)
- **Multi-mode Energy Management for Multi-tier Server Clusters** 270
Tibor Horvath, Kevin Skadron (*University of Virginia*)

Session 9: Programming the Memory Hierarchy

- **A Tuning Framework for Software-Managed Memory Hierarchies** 280
Manman Ren, Ji Young Park, Mike Houston, Alex Aiken, William J. Dally (*Stanford University*)
 - **Hybrid Access-Specific Software Cache Techniques for the Cell BE Architecture** 292
Marc González (*Universitat Politècnica de Catalunya and Barcelona Supercomputing Center*),
Nikola Vujic (*Barcelona Supercomputing Center*),
Xavier Martorell, Eduard Ayguadé (*Universitat Politècnica de Catalunya & Barcelona Supercomputing Center*),
Alexandre E. Eichenberger, Tong Chen, Zehra Sura, Tao Zhang, Kevin O'Brien,
Kathryn O'Brien (*IBM T.J. Watson Research Center*)
 - **COMIC: A Coherent Shared Memory Interface for Cell BE** 303
Jaejin Lee, Sangmin Seo, Chihun Kim, Junghyun Kim, Posung Chun (*Seoul National University*),
Zehra Sura (*IBM T. J. Watson Research Center*), Jungwon Kim, SangYong Han (*Seoul National University*)
- Author Index** 315