

2013 22nd International Conference on Parallel Architectures and Compilation Techniques

(PACT 2013)

**Edinburgh, United Kingdom
7-11 September 2013**



IEEE Catalog Number: CFP13073-POD
ISBN: 978-1-4799-1020-5

Table of Contents

PACT 2013 Organization	zk
-------------------------------------	----

PACT 2013 Sponsors and Supporters	zk
--	----

Keynote Addresses

- **Keynote: A Comprehensive Approach to HW/SW Codesign** 1
David Kuck (*Intel Corporation*)
- **Keynote: Parallel Programming for Mobile Computing** 3
Chin Ca?caval (*Qualcomm Research Silicon Valley*)
- **Keynote: Towards Automatic Resource Management in Parallel Architectures** 5
Per Stenström (*Chalmers University of Technology*)

Session 1A: Compilers

- **INSPIRE: The Insieme Parallel Intermediate Representation** 7
Herbert Jordan (*University of Innsbruck*),
Simone Pellegrini, Peter Thoman, Klaus Kofler, Thomas Fahringer (*Institute of Computer Science/University of Innsbruck*)
- **Parallel Flow-Sensitive Pointer Analysis by Graph-Rewriting** 19
Vaivaswatha Nagaraj, R Govindarajan (*Indian Institute of Science*)
- **Interprocedural Strength Reduction of Critical Sections in Explicitly-Parallel Programs** 29
Rajkishore Barik (*Intel Labs*), Jisheng Zhao, Vivek Sarkar (*Rice University*)

Session 1B: Power & Energy

- **ThermOS: System Support for Dynamic Thermal Management of Chip Multi-Processors** 41
Filippo Sironi (*Politecnico di Milano*), Martina Maggio (*Lund University*),
Riccardo Cattaneo, Giovanni F. Del Nero, Donatella Sciuto, Marco D. Santambrogio (*Politecnico di Milano*)
- **Coordinated Power-Performance Optimization in Manycores** 51
Hiroshi Sasaki, Satoshi Imamura, Koji Inoue (*Kyushu University*)
- **An Opportunistic Prediction-Based Thread Scheduling to Maximize Throughput/Watt in AMPs** 63
Arunachalam Annamalai, Rance Rodrigues, Israel Koren, Sandip Kundu (*University of Massachusetts at Amherst*)

Session 2A: GPU & Energy

- **APOGEE: Adaptive Prefetching on GPUs for Energy Efficiency** 73
Ankit Sethia (*University of Michigan*), Ganesh Dasika (*ARM R&D*),
Mehrzaad Samadi, Scott Mahlke (*University of Michigan*)
- **Parallel Frame Rendering: Trading Responsiveness for Energy on a Mobile GPU** 83
Jose-Maria Arnau, Joan-Manuel Parcerisa (*Universitat Politecnica de Catalunya*),
Polychronis Xekalakis (*Intel Corporation*)
- **Exploring Hybrid Memory for GPU Energy Efficiency Through Software-Hardware Co-Design** 93
Bin Wang (*Auburn University*), Bo Wu (*College of William & Mary*), Dong Li (*Oak Ridge National Laboratory*),
Xipeng Shen (*College of William & Mary*), Weikuan Yu, Yizheng Jiao (*Auburn University*),
Jeffrey S. Vetter (*Oak Ridge National Laboratory*)

Session 2B: Memory System Management

- **S-CAVE: Effective SSD Caching to Improve Virtual Machine Storage Performance** 103
Tian Luo, Siyuan Ma, Rubao Lee, Xiaodong Zhang (*The Ohio State University*), Deng Liu (*VMware Inc.*),
Li Zhou (*Facebook Inc.*)

- **Writeback-Aware Bandwidth Partitioning for Multi-Core Systems with PCM** 113
Miao Zhou, Yu Du, Bruce R. Childers, Rami Melhem, Daniel Mossé (*University of Pittsburgh*)
- **L1-Bandwidth Aware Thread Allocation in Multicore SMT Processors** 123
Josué Feliu, Julio Sahuquillo, Salvador Petit, José Duato (*Universitat Politècnica de València*)

Session 3: Best Papers

- **A Unified View of Non-Monotonic Core Selection and Application Steering in Heterogeneous Chip Multiprocessors** 133
Sandeep Navada, Niket K. Choudhary, Salil V. Wadhavkar (*Qualcomm*),
Eric Rotenberg (*North Carolina State University*)
- **Memory-Centric System Interconnect Design with Hybrid Memory Cubes** 145
Gwangsun Kim, John Kim (*Korea Advanced Institute of Science and Technology*),
Jung Ho Ahn, Jaeha Kim (*Seoul National University*)
- **Neither More Nor Less: Optimizing Thread-Level Parallelism for GPGPUs** 157
Onur Kay?ran, Adwait Jog, Mahmut T. Kandemir, Chita R. Das (*The Pennsylvania State University*)
- **SMT-Centric Power-Aware Thread Placement in Chip Multiprocessors** 167
Augusto Vega, Alper Buyuktosunoglu, Pradip Bose (*IBM T. J. Watson Research Center*)

Session 4A: Runtime & Scheduling

- **Fairness-Aware Scheduling on Single-ISA Heterogeneous Multi-Cores** 177
Kenzo Van Craeynest (*Ghent University & ExaScience Lab*), Shoaib Akram (*Ghent University*),
Wim Heirman (*Ghent University & ExaScience Lab*), Aamer Jaleel (*Intel Corporation*),
Lieven Eeckhout (*Ghent University*)
- **DANBI: Dynamic Scheduling of Irregular Stream Programs for Many-Core Systems** 189
Changwoo Min (*Sungkyunkwan University and Samsung Electronics*), Young Ik Eom (*Sungkyunkwan University*)
- **An Empirical Model for Predicting Cross-Core Performance Interference on Multicore Processors** 201
Jiacheng Zhao (*Institute of Computing Technology, CAS & University of Chinese Academy of Sciences*),
Huimin Cui (*Chinese Academy of Science*), Jingling Xue (*University of New South Wales*),
Xiaobing Feng (*Institute of Computing Technology, CAS*),
Youliang Yan, Wensen Yang (*Huawei Technologies Co., Ltd.*)

Session 4B: Caches & Memory Hierarchy (1)

- **Jigsaw: Scalable Software-Defined Caches** 213
Nathan Beckmann, Daniel Sanchez (*Massachusetts Institute of Technology*)
- **Managing Shared Last-Level Cache in a Heterogeneous Multicore Processor** 225
Vineeth Mekkat, Anup Holey, Pen-Chung Yew, Antonia Zhai (*University of Minnesota*)
- **Reshaping Cache Misses to Improve Row-Buffer Locality in Multicore Systems** 235
Wei Ding, Jun Liu, Mahmut Kandemir, Mary Jane Irwin (*The Pennsylvania State University*)

Session 5A: GPU

- **Transparent CPU-GPU Collaboration for Data-Parallel Kernels on Heterogeneous Systems** 245
Janghaeng Lee, Mehrzad Samadi, Yongjun Park, Scott Mahlke (*University of Michigan*)
- **Starchart: Hardware and Software Optimization Using Recursive Partitioning Regression Trees** 257
Wenhao Jia (*Princeton University*), Kelly A. Shaw (*University of Richmond*),
Margaret Martonosi (*Princeton University*)
- **RSVM: A Region-Based Software Virtual Memory for GPU** 269
Feng Ji (*NC State University*), Heshan Lin (*Virginia Tech*), Xiaosong Ma (*NC State University & ORNL*)

Session 5B: Caches & Memory Hierarchy (2)

- **The Case for a Scalable Coherence Protocol for Complex On-Chip Cache Hierarchies in Many-Core Systems** 279
Lucia G. Menezes, Valentin Puente, Jose Angel Gregorio (*University of Cantabria*)
- **Meeting Midway: Improving CMP Performance with Memory-Side Prefetching** 289
Praveen Yedlapalli, Jagadish Kotra, Emre Kultursay, Mahmut Kandemir, Chita R. Das, Anand Sivasubramaniam (*The Pennsylvania State University*)
- **Building Expressive, Area-Efficient Coherence Directories** 299
Lei Fang, Peng Liu, Qi Hu (*Zhejiang University*), Michael C. Huang (*University of Rochester*), Guofan Jiang (*IBM China Systems and Technology Lab*)

Session 6A: Network, Debugging & Microarchitecture

- **Traffic Steering Between a Low-Latency Unswitched TL Ring and a High-Throughput Switched On-Chip Interconnect** 309
Jungju Oh, Alenka Zajic, Milos Prvulovic (*Georgia Institute of Technology*)
- **McRouter: Multicast Within a Router for High Performance Network-on-Chips** 319
Yuan He (*The University of Tokyo*), Hiroshi Sasaki (*Kyushu University*), Shinobu Miwa, Hiroshi Nakamura (*The University of Tokyo*)
- **Concurrent Predicates: A Debugging Technique for Every Parallel Programmer** 331
Justin Gottschlich, Gilles Pokam, Cristiano Pereira, Youfeng Wu (*Intel Corporation*)
- **Breaking SIMD Shackles with an Exposed Flexible Microarchitecture and the Access Execute PDG** 341
Venkatraman Govindaraju, Tony Nowatzki, Karthikeyan Sankaralingam (*University of Wisconsin-Madison*)

Session 6B: Compiler Optimization

- **Vectorization Past Dependent Branches Through Speculation** 353
Majedul Haque Sujon (*University of Texas at San Antonio*), R. Clint Whaley (*Louisiana State University*), Qing Yi (*University of Colorado Colorado Springs*)
- **Automatic Vectorization of Tree Traversals** 363
Youngjoon Jo, Michael Goldfarb, Milind Kulkarni (*Purdue University*)
- **Generating Efficient Data Movement Code for Heterogeneous Architectures with Distributed-Memory** 375
Roshan Dathathri, Chandan Reddy, Thejas Ramashekar, Uday Bondhugula (*Indian Institute of Science*)
- **Automatic OpenCL Work-Group Size Selection for Multicore CPUs** 387
Sangmin Seo (*ManyCoreSoft*), Jun Lee, Gangwon Jo, Jaejin Lee (*Seoul National University*)

Posters

- **TCPT - Thread Criticality-driven Prefetcher Throttling** 399
Biswabandan Panda, Shankar Balachandran (*IIT Madras*)
- **Do Inputs Matter? Using Data-Dependence Profiling to Evaluate Thread Level Speculation in BG/Q** 401
Arnaboy Bhattacharyya (*University of Alberta*)
- **Can Lock-Free and Combining Techniques Co-Exist? A Novel Approach on Concurrent Queue** 403
Changwoo Min, Young Ik Eom (*Sungkyunkwan University*)
- **Task Sampling: Computer Architecture Simulation in the Many-Core Era** 405
Thomas Grass (*Universitat Politècnica de Catalunya (BarcelonaTech) & Barcelona Supercomputing Center*)
- **PS-Cache: An Energy-Efficient Cache Design for Chip Multiprocessors** 407
Joan J. Valls (*Universitat Politècnica de València*), Alberto Ros (*Universidad de Murcia*), Julio Sahuquillo, María Engracia Gómez (*Universitat Politècnica de València*)

• Dynamic Memory Access Monitoring Based on Tagged Memory	409
Mikhail Gorelov (<i>Moscow Institute of Physics and Technology/MCST</i>), Lev Mukhanov (<i>MCST</i>)	
• Exposing ILP in Custom Hardware with a Dataflow Compiler IR	411
Ali Mustafa Zaidi (<i>University of Cambridge</i>)	
Author Index	412