

2009 42nd Annual IEEE/ACM International Symposium on Microarchitecture

(MICRO 2009)

**New York, New York, USA
12 – 16 December 2009**



**IEEE Catalog Number: CFP09071-PRT
ISBN: 978-1-4244-6312-1**

Table of Contents

Keynote 1

- **POWER7 Multi-core Processor Design** 1
Balaram Sinharoy (*IBM*)

Session 1: Memory System

Session Chair: Engin İpek (*University of Rochester*)

- **Characterizing and Mitigating the Impact of Process Variations on Phase Change based Memory Systems** 2
Wangyuan Zhang, Tao Li (*University of Florida*)
- **Enhancing Lifetime and Security of PCM-Based Main Memory with Start-Gap Wear Leveling** 14
Moinuddin K. Qureshi, John Karidis, Vijayalakshmi Srinivasan, Michele Franceschini, Luis Lastras, Bulent Abali (*IBM Research*)
- **Characterizing Flash Memory: Anomalies, Observations, and Applications** 24
Laura M. Grupp, Adrian M. Caulfield, Joel Coburn, Steven Swanson, Eitan Yaakobi, Paul H. Siegel, Jack K. Wolf (*University of California, San Diego*)
- **Complexity Effective Memory Access Scheduling for Many-Core Accelerator Architectures** 34
George L. Yuan, Ali Bakhoda, Tor M. Aamodt (*University of British Columbia*)

Session 2A: Optimization

Session Chair: Kim Hazelwood (*University of Virginia*)

- **Qilin: Exploiting Parallelism on Heterogeneous Multiprocessors with Adaptive Mapping** 45
Chi-Keung Luk (*Intel*) Sunpyo Hong, Hyesoon Kim (*Georgia Tech*)
- **DDT: Design and Evaluation of a Dynamic Program Analysis for Optimizing Data Structure Usage** 56
Changhee Jung, Nathan Clark (*Georgia Institute of Technology*)
- **Tree Register Allocation** 67
Hongbo Rong (*Microsoft Corp.*)
- **Portable Compiler Optimization Across Embedded Programs and Microarchitectures using Machine Learning** 78
Christophe Dubach, Timothy M. Jones, Edwin V. Bonilla (*University of Edinburgh*)
Grigori Fursin (*INRIA Saclay, France*) Michael F.P. O'Boyle (*University of Edinburgh*)

Session 2B: Fault Tolerance

Session Chair: Ed Suh (*Cornell University*)

- **Improving Cache Lifetime Reliability at Ultra-low Voltages** 89
Zeshan Chishti, Alaa R. Alameldeen, Chris Wilkerson, Wei Wu, Shih-Lien Lu (*Intel*)
- **ZerehCache: Armoring Cache Architectures in High Defect Density Technologies** 100
Amin Ansari, Shantanu Gupta, Shuguang Feng, Scott Mahlke (*University of Michigan, Ann Arbor*)
- **Low Vccmin Fault-Tolerant Cache with Highly Predictable Performance** 111
Jaume Abellà, Javier Carretero, Pedro Chaparro, Xavier Vera (*Intel Barcelona Research Center*)
Antonio González, (*Intel Barcelona Research Center, Universitat Politècnica de Catalunya*)
- **mSWAT: Low-Cost Hardware Fault Detection and Diagnosis for Multicore Systems** 122
Siva Kumar Sastry Hari, Man-Lap Li, Pradeep Ramachandran, Byn Choi, Sarita V. Adve (*University of Illinois, Urbana-Champaign*)

Session 3A: Transactional Memory/Memory Consistency

Session Chair: Milo M. K. Martin (*University of Pennsylvania*)

- **BulkCompiler: High-Performance Sequential Consistency through Cooperative Compiler and Hardware Support** 133
Wonsun Ahn, Shanxiang Qi, Marios Nicolaides, Josep Torrellas (*University of Illinois, Urbana-Champaign*)
Jae-Woo Lee, Xing Fang, Samuel Midkiff (*Purdue University*) David Wong (*Intel*)
- **EazyHTM: Eager-Lazy Hardware Transactional Memory** 145
Saša Tomić, Cristian Perfumo (*BSC-Microsoft Research Centre, Universitat Politècnica de Catalunya*)
Chinmay Kulkarni, (*BSC-Microsoft Research Centre, BITS Pilani*)
Adrià Armejach, (*BSC-Microsoft Research Centre, Universitat Politècnica de Catalunya*)
Adrián Cristal, Osman Unsal (*BSC-Microsoft Research Centre*)
Tim Harris (*Microsoft Research Cambridge*) Mateo Valero (*BSC-Microsoft Research Centre*)
- **Proactive Transaction Scheduling for Contention Management** 156
Geoffrey Blake, Ronald G. Dreslinski, Trevor Mudge (*University of Michigan, Ann Arbor*)

Session 3B: Power

Session Chair: David Brooks (*Harvard University*)

- **Into the Wild: Studying Real User Activity Patterns to Guide Power Optimization for Mobile Architectures** 168
Alex Shye, Benjamin Scholbrock, Gokhan Memik (*Northwestern University*)
- **A Microarchitecture-based Framework for Pre- and Post-Silicon Power Delivery Analysis** 179
Mahesh Ketkar, Eli Chiprout (*Intel*)
- **Reducing Peak Power with a Table-Driven Adaptive Processor Core** 189
Vasileios Kontorinis, Amirali Shayan (*University of California, San Diego*)
Rakesh Kumar (*University of Illinois, Urbana-Champaign*)
Dean Tullsen (*University of California, San Diego*)

Session 4A: 3D/Caches

Session Chair: Trevor Mudge (*University of Michigan – Ann Arbor*)

- **Extending the Effectiveness of 3D-Stacked DRAM Caches with an Adaptive Multi-Queue Policy** 201
Gabriel H. Loh (*Georgia Institute of Technology*)
- **An Hybrid eDRAM/SRAM Macrocell to Implement First-level Data Caches** 213
Alejandro Valero, Julio Sahuquillo, Salvador Petit, Vicente Lorente (*Universidad Politècnica de Valencia*)
Ramón Canal (*Universitat Politècnica de Catalunya*)
Pedro López, José Duato (*Universidad Politècnica de Valencia*)
- **Variation-Tolerant Non-Uniform 3D Cache Management in Die Stacked Multicore Processor** 222
Bo Zhao, Yu Du, Youtao Zhang, Jun Yang (*University of Pittsburgh*)

Session 4B: On-chip Networks

Session Chair: Manish Vachharajani (*University of Colorado, Boulder*)

- **In-Network Coherence Filtering: Snoopy Coherence without Broadcasts** 232
Niket Agarwal, Li-Shiuan Peh, Niraj K. Jha (*Princeton University*)
- **SCARAB: A Single Cycle Adaptive Routing and Bufferless Network** 244
Mitchell Hayenga (*University of Wisconsin, Madison*) Natalie Enright-Jerger (*University of Toronto*)
Mikko Lipasti (*University of Wisconsin, Madison*)
- **Low-Cost Router Microarchitecture for On-Chip Networks** 255
John Kim (*KAIST*)

Keynote 2

- **Why Design Must Change: Rethinking Digital Design**267
Mark Horowitz (*Stanford University*)

Session 5: On-chip Networks

Session Chair: Tom Conte (*Georgia Tech*)

- **Preemptive Virtual Clock: A Flexible, Efficient and Cost-effective QOS Scheme for Networks-on-Chip**268
Boris Grot, Stephen W. Keckler (*The University of Texas, Austin*) Onur Mutlu (*Carnegie Mellon University*)
- **Application-Aware Prioritization Mechanisms for On-Chip Networks**280
Reetuparna Das (*Pennsylvania State University*) Onur Mutlu (*Carnegie Mellon University*)
Thomas Moscibroda (*Microsoft Research*) Chita R. Das (*Pennsylvania State University*)
- **A Case for Dynamic Frequency Tuning in On-Chip Networks**292
Asit K. Mishra, Reetuparna Das, Soumya Eachempati (*Pennsylvania State University*)
Ravi Iyer (*Intel*) N. Vijaykrishnan, Chita R. Das (*Pennsylvania State University*)
- **Light Speed Arbitration and Flow Control for Nanophotonic Interconnects**304
Dana Vantrease (*University of Wisconsin, Madison*) Nathan Binkert, Robert Schreiber (*HP Labs*)
Mikko Lipasti (*University of Wisconsin, Madison*)

Session 6A: Memory Systems

Session Chair: Thomas F. Wenisch (*University of Michigan, Ann Arbor*)

- **Coordinated Control of Multiple Prefetchers in Multi-Core Systems**316
Eiman Ebrahimi (*The University of Texas, Austin*) Onur Mutlu (*Carnegie Mellon University*)
Chang Joo Lee, Yale N. Patt (*The University of Texas, Austin*)
- **Improving Memory Bank-Level Parallelism in the Presence of Prefetching**327
Chang Joo Lee, Veynu Narasiman (*The University of Texas, Austin*)
Onur Mutlu (*Carnegie Mellon University*) Yale N. Patt (*The University of Texas, Austin*)
- **ESKIMO - Energy Savings using Semantic Knowledge of Inconsequential Memory Occupancy for DRAM subsystem**337
Ciji Isen, Lizy John (*University of Texas, Austin*)
- **Flip-N-Write: A Simple Deterministic Technique to Improve PRAM Write Performance, Energy and Endurance**347
Sangyeun Cho, Hyunjin Lee (*University of Pittsburgh*)

Session 6B: Reconfigurability/Security

Session Chair: Martha Kim (*Columbia University*)

- **Using a Configurable Processor Generator for Computer Architecture Prototyping**358
Alex Solomatnikov, Amin Firoozshahian (*Hicamp Systems Inc.*) Ofer Shacham, Zain Asgar,
Megan Wachs, Wajahat Qadeer, Stephen Richardson, Mark Horowitz (*Stanford University*)
- **Polymorphic Pipeline Array: A Flexible Multicore Accelerator with Virtualized Execution for Mobile Multimedia Applications**370
Hyunchul Park, Yongjun Park, Scott Mahlke (*University of Michigan, Ann Arbor*)
- **Ordering Decoupled Metadata Accesses in Multiprocessors**381
Hari Kannan (*Stanford University*)
- **Control Flow Obfuscation with Information Flow Tracking**391
Haibo Chen, Liwei Yuan, Xi Wu, Binyu Zang (*Fudan University*)
Bo Huang (*Intel*) Pen-chung Yew (*University of Minnesota*)

Session 7A: Caches

Session Chair: Karin Strauss (*Microsoft*)

- **Pseudo-LIFO: The Foundation of a New Family of Replacement Policies for Last-level Caches**401
Mainak Chaudhuri (*Indian Institute of Technology, Kanpur*)
- **Comparing Cache Architectures and Coherency Protocols on x86-64 Multicore SMP Systems**413
Daniel Hackenberg, Daniel Molka, Wolfgang E. Nagel (*ZIH, TU Dresden*)
- **A Tagless Coherence Directory**423
Jason Zebchuk (*University of Toronto*) Vijayalakshmi Srinivasan, Moinuddin K. Qureshi (*IBM*)
Andreas Moshovos (*University of Toronto*)

Session 7B: Variations/Power

Session Chair: Hyesoon Kim (*Georgia Tech*)

- **Tribeca: Design for PVT Variations with Local Recovery and Fine-grained Adaptation**435
Meeta Sharma Gupta (*Harvard University*) Jude A. Rivers, Pradip Bose (*IBM*)
Gu-Yeon Wei, David Brooks (*Harvard University*)
- **The BubbleWrap Many-core: Popping Cores for Sequential Acceleration**447
Ulya R. Karpuzcu, Brian Greskamp, Josep Torrellas (*University of Illinois, Urbana-Champaign*)
- **Multiple Clock and Voltage Domains for Chip Multi Processors**459
Efraim Rotem (*Intel*) Ran Ginosar (*Technion*) Avi Mendelson (*Microsoft*) Uri Weiser (*Technion*)

Session 8A: Potpourri

Session Chair: Anne Bracy (*Intel*)

- **McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures**469
Sheng Li (*University of Notre Dame*) Jung Ho Ahn (*HP Labs*)
Richard D. Strong (*University of California, San Diego*) Jay B. Brockman (*University of Notre Dame*)
Dean M. Tullsen (*University of California, San Diego*) Norman P. Jouppi (*HP Labs*)
- **Characterizing the Resource-Sharing Levels in the UltraSPARC T2 Processor**481
Vladimir Ćakarević, Petar Radojković (*Barcelona Supercomputing Center*)
Javier Verdú, Alex Pajuelo (*Universitat Politècnica de Catalunya*)
Francisco J. Cazorla (*Barcelona Supercomputing Center*)
Mario Nemirovsky (*Barcelona Supercomputing Center, ICREA*)
Mateo Valero (*Barcelona Supercomputing Center, Universitat Politècnica de Catalunya*)
- **Execution Leases: A Hardware-Supported Mechanism for Enforcing Strong Non-Interference**493
Mohit Tiwari, Xun Li, Hassan M G Wassel, Frederic T. Chong,
Timothy Sherwood (*University of California, Santa Barbara*)

Session 8B: Caches

Session Chair: Amir Roth (*University of Pennsylvania*)

- **Optimizing Shared Cache Behavior of Chip Multiprocessors**505
Mahmut Kandemir, Sai Prasanth Muralidhara, Sri Hari Krishna Narayanan,
Yuanrui Zhang (*Pennsylvania State University*) Ozcan Ozturk (*Bilkent University*)
- **SHARP Control: Controlled Shared Cache Management in Chip Multiprocessors**517
Shekhar Srikantaiah, Mahmut Kandemir, Qian Wang (*Pennsylvania State University*)
- **Adaptive Line Placement with the Set Balancing Cache**529
Dyer Rolan, Basilio B. Fraguera, Ramón Doallo (*Universidade da Coruña*)

Session 9: Debugging

Session Chair: Emmett Witchel (*University of Texas, Austin*)

- **Light64: Lightweight Hardware Support for Race Detection during Systematic Testing of Parallel Programs**541
Adrian Nistor, Darko Marinov, Josep Torrellas (*University of Illinois, Urbana-Champaign*)
- **Finding Concurrency Bugs with Context-Aware Communication Graphs**553
Brandon Lucia, Luis Ceze (*University of Washington*)
- **Offline Symbolic Analysis for Multi-Processor Execution Replay**.....564
Dongyoon Lee (*University of Michigan, Ann Arbor*) Mahmoud Said (*Western Michigan University*)
Satish Narayanasamy (*University of Michigan, Ann Arbor*) Zijiang (James) Yang (*Western Michigan University*)
Cristiano Pereira (*Intel*)
- **Architecting a Chunk-based Memory Race Recorder in Modern CMPs**.....576
Gilles Pokam, Cristiano Pereira, Klaus Danne, Rolf Kassa, Ali-Reza Adl-Tabatabai (*Intel*)

Author Index