

Proceedings

*The Fifteenth International Symposium on
High-Performance Computer Architecture*

HPCA-15 2009

February 14-18, 2009 • Raleigh, NC

Sponsored by

IEEE Computer Society Technical Committee on Computer Architecture



Los Alamitos, California
Washington • Brussels • Tokyo

Table of Contents

Message from the General Chairs
Message from the Program Chair
Organizing Committee/Steering Committee
Program Committee
External Reviewers

Keynote I

Chair: Tom Conte, Georgia Tech

An Intelligent IT Infrastructure for the Future..... 3
Prith Banerjee, HP Labs

Session 1 – Best Paper Nominees

Chair: Josep Torrellas, UIUC

Techniques for Bandwidth-Efficient Prefetching of Linked Data
Structures in Hybrid Prefetching Systems 7
Eiman Ebrahimi (University of Texas at Austin), Onur Mutlu (Carnegie Mellon), Yale Patt (University of Texas at Austin)

Voltage Emergency Prediction: Using Signatures to Reduce Operating Margins..... 18
Vijay Janapa Reddi, Meeta Gupta, Glenn Holloway, Gu Yeon Wei, Michael D. Smith, David Brooks (Harvard University)

A Low-Radix and Low-Diameter 3D Interconnection Network Design 30
Yi Xu, Yu Du, Bo Zhao, Xiuyi Zhou, Youtao Zhang, Jun Yang (University of Pittsburgh)

Session 2A – Multicore Cache Architectures

Chair: Gabriel Loh, Georgia Tech.

Adaptive Spill-Receive for Robust High-Performance Caching in CMPs 45
Moinuddin K. Qureshi (IBM Research)

Design and Implementation of Software-Managed Caches for Multicores with Local Memory 55
Sangmin Seo, Jaejin Lee (Seoul National University), Zehra Sura (IBM Research)

In-Network Snoop Ordering (INSO): Snoopy Coherence on Unordered Interconnects 67
Niket Agarwal, Li-Shiuan Peh, Niraj Jha (Princeton University)

Practical Off-chip Meta-data for Temporal Memory Streaming 79
Thomas Wenisch (University of Michigan), Michael Ferdman, Anastasia Ailamaki (Carnegie Mellon University), Babak Falsafi (EPFL), Andreas Moshovos (University of Toronto)

Session 2B – Reliability

Chair: Todd Austin, Univ. of Michigan

Soft Error Vulnerability Aware Process Variation Mitigation 93
Xin Fu, Tao Li, Jose Fortes (University of Florida)

Accurate Microarchitecture-Level Fault Modeling for Studying Hardware Faults 105
Manlap Li, Pradeep Ramachandran, R. Ulya Karpuzcu, Siva Kumar Sastry Hari, Sarita V. Adve, Vikram S. Adve (University of Illinois, Urbana-Champaign)

Eliminating Microarchitectural Dependency from Architectural Vulnerability 117
Vilas Sridharan, David R. Kaeli (Northeastern University)

Versatile Prediction and Fast Estimation of Architectural Vulnerability Factor
from Processor Performance Metrics 129
Lide Duan, Bin Li, Lu Peng (Louisiana State University)

Panel (joint with PPOP)

Opportunities Beyond Single-Core Microprocessors 143
Moderator: Mark Hill, Univ. of Wisconsin, Madison
Panelists: Sarita V. Adve (Univ. of Illinois), David A. Bader (Georgia Tech), William Dally (Stanford Univ.), William Harrod (DARPA), Vivek Sarkar (Rice Univ.)

Keynote II (joint with PPOP)

Chair: Yan Solihin, NC State Univ.

Multi-core Demands Multi-interfaces 147
Yale Patt, U.T. Austin

Session 3A – On-Chip Networks-I

Chair: Timothy M. Pinkston, USC

Elastic Buffer Flow Control for On-Chip Networks 151
George Micheliogiannakis, James Balfour, William J. Dally (Stanford University)

Express Cube Topologies for On-Chip Interconnects 163
Boris Grot, Joel Hestness (University of Texas at Austin), Stephen W. Keckler (University of Texas at Austin), Onur Mutlu (Carnegie Mellon)

Design and Evaluation of a Hierarchical On-Chip Interconnect for Next-Generation CMPs 175
Reetuparna Das, Soumya Eachempati, Asit K Mishra, Vijaykrishnan Narayanan, Chita Das (Pennsylvania State University)

Session 3B – Processor Microarchitecture-I

Chair: Pradip Bose, IBM T. J. Watson Research

Architectural Contesting 189
Hashem Hashemi Najaf-abadi, Eric Rotenberg (NC State University)

Lightweight Predication Support for Out of Order Processors 201
Mark Stephenson, Lixin Zhang, Ram Rangan (IBM Research)

BlueShift: Designing Processors for Timing Speculation from the Ground Up 213
Brian Greskamp, Lu Wan, Ulya R. Karpuzcu, Jeffrey J. Cook, Josep Torrellas, Deming Chen, Craig Zilles (University of Illinois)

Session 4A – NUCA and 3-D Stacked Memory Hierarchies

Chair: Doug Burger, Microsoft Research

PageNUCA: Selected Policies for Page-grain Locality Management in Large Shared Chip-multiprocessor Caches	227
<i>Mainak Chaudhuri (Indian Institute of Technology, Kanpur)</i>	
A Novel Architecture of the 3D Stacked MRAM L2 Cache for CMPs.....	239
<i>Guangyu Sun, Xiangyu Dong, Yuan Xie (Penn State University), Jian Li (IBM), Yiran Chen (Seagate Tech)</i>	
Dynamic Hardware-Assisted Software-Controlled Page Placement to Manage Capacity Allocation and Sharing within Large Caches	250
<i>Manu Awasthi, Kshitij Sudan, Rajeev Balasubramonian, John Carter (University of Utah)</i>	
Optimizing Communication and Capacity in a 3D Stacked Reconfigurable Cache Hierarchy	262
<i>Niti Madan (University of Utah), Li Zhao (Intel), Naveen Muralimanohar, Aniruddha Udipi, Rajeev Balasubramonian (University of Utah), Ravishankar Iyer, Srihari Makineni, Donald Newell (Intel)</i>	

Session 4B – Power/Performance-Efficient Architectures and Accelerators

Chair: Chair: David Brooks, Harvard University

Reconciling Specialization and Flexibility Through Compound Circuits	277
<i>Sami Yehia, Sylvain Girbal (Thales Research & Technology), Hugues Berry, Olivier Temam (INRIA),</i>	
CAMP: A Technique to Estimate Per-Structure Power at Run-time using a Few Simple Parameters. 289	
<i>Michael D. Powell, Arijit Biswas, Joel Emer, Shubhendu S. Mukherjee (Intel), Basit R. Sheikh (Cornell), Shrirang Yardi (Virginia Tech)</i>	
Variation-Aware Dynamic Voltage/Frequency Scaling	301
<i>Sebastian Herbert, Diana Marculescu (Carnegie Mellon University)</i>	
Bridging the Computation Gap Between Programmable Processors and Hardwired Accelerators.....	313
<i>Kevin Fan, Manjunath Kudlur, Ganesh Dasika, Scott Mahlke (University of Michigan)</i>	
Industrial Perspectives Panel (joint with PPOPP)	325
<i>Moderator: Partha Ranganathan (HP Labs)</i>	
<i>Panelists: Rick Hetherington (Sun), David Luebke (Nvidia), Qureshi Moinuddin (IBM T. J. Watson Research), Daniel Reed (Microsoft Research)</i>	

Session 5A – Performance Modeling and Analysis

Chair: Olivier Temam, INRIA

A First-Order Fine-Grained Multithreaded Throughput Model.....	329
<i>Xi E. Chen, Tor M. Aamodt (University of British Columbia)</i>	
Characterization of Direct Cache Access on Multi-core Systems and 10GbE	341
<i>Amit Kumar, Ram Huggahalli, Srihari Makineni (Intel)</i>	

Session 5B – On-Chip Networks-II

Chair: Jun Yang, Univ. of Pittsburgh

MRR: Enabling Fully Adaptive Multicast routing for CMP interconnection networks 355
Pablo Abad Fidalgo, Valentin Puente Varona, Jose Angel Gregorio Monasterio (University of Cantabria)

Prediction Router: Yet Another Low Latency On-Chip Router Architecture 367
*Hiroki Matsutan (Keio University), Michihiro Koibuchi (National Institute of Informatics),
Hideharu Amano (Keio University), Tsutomu Yoshinaga (The University of Electro-Communications)*

Session 6A – Security, Verification, and Validation

Chair: Youtao Zhang, Univ. of Pittsburgh

Fast Complete Memory Consistency Verification 381
*Yunji Chen, Yi Lv (Chinese Academy of Sciences), Tianshi Chen (University of Science and Technology of
China), Haihua Shen, Pengyu Wang, Hong Pan, Weiwu Hu (Chinese Academy of Sciences)*

Hardware-Software Integrated Approaches to Defend Against Software
Cache-based Side Channel Attacks..... 393
*Jingfei Kong (University of Central Florida), Onur Aciicmez, Jean-Pierre Seifert
(Samsung Information Systems America), Huiyang Zhou (University of Central Florida)*

DACOTA: Post-silicon validation of the memory subsystem in multi-core designs 405
Andrew DeOrio, Ilya Wagner, Valeria Bertacco (University of Michigan)

Session 6B – Processor Microarchitecture-II

Chair: Eric Rotenberg, NC State Univ.

Criticality-Based Optimizations for Efficient Load Processing..... 419
*Samantika Subramaniam (Georgia Institute of Technology), Anne C. Bracy, Hong Wang (Intel),
Gabriel H. Loh (Georgia Institute of Technology)*

iCFP: Tolerating All Level Cache Misses in In-Order Processors 431
Andrew Hilton, Santosh Nagarakatte, Amir Roth (University of Pennsylvania)

Feedback Mechanisms for Improving Probabilistic Memory Prefetching 443
Ibrahim Hur (IBM), Calvin Lin (University of Texas at Austin)

Keynote III (joint with PPOP)

Chair: Keshav Pingali, Univ. of Texas, Austin

How to build Programmable Multi-Core Chips..... 457
Jack Dennis, MIT

Author Index