

Computer Architecture News




*A Publication of the
Association for Computing Machinery
Special Interest Group on Computer Architecture*

Vol. 34, No. 2 - May 2006

ISCA '06 PROCEEDINGS

***PROCEEDINGS OF THE 33RD
ANNUAL INTERNATIONAL
SYMPOSIUM ON COMPUTER
ARCHITECTURE***

June 17-21, 2006 Boston, Massachusetts



Proceedings

33rd International Symposium on Computer Architecture

17-21 June 2006 • Boston, Massachusetts


Sponsored by

IEEE Computer Society TCCA
ACM SIGARCH



Los Alamitos, California

Washington • Tokyo



Proceedings

ISCA 2006

Table of Contents

Message from the General Chair	x
Message from the Program Chair	xi
Officers	xii
Reviewers	xiv
SIGARCH Guidelines	xvii

Keynote 1

Chair: Mateo Valero, *Universitat Politècnica de Catalunya*

Computer Architecture Research and Future Microprocessors: Where Do We Go from Here?	2
<i>Yale Patt, University of Texas at Austin</i>	

Session 1: Interconnection Networks

Chair: Wen-Mei W. Hwu, *University of Illinois at Urbana-Champaign*

A Gracefully Degrading and Energy-Efficient Modular Router Architecture For On-Chip Networks	4
<i>Jongman Kim, Chrysostomos Nicopoulos, Dongkook Park, Vijaykrishnan Narayanan, Mazin S. Yousif, and Chita R. Das</i>	
The BlackWidow High-Radix Clos Network.....	16
<i>Steve Scott, Dennis Abts, John Kim, and William J. Dally</i>	

Session 2: Memory Models

Chair: José F. Martínez, *Cornell University*

Memory Model = Instruction Reordering + Store Atomicity	29
<i>Arvind and Jan-Willem Maessen</i>	
Conditional Memory Ordering	41
<i>Christoph von Praun, Harold W. Cain, Jong-Deok Choi, and Kyung Dong Ryu</i>	
Architectural Semantics for Practical Transactional Memory	53
<i>Austen McDonald, JaeWoong Chung, Brian D. Carlstrom, Chi Cao Minh, Hassan Chafi, Christos Kozyrakis, and Kunle Olukotun</i>	

Session 3: Power and Thermal Management

Chair: Josep Torrellas, *University of Illinois at Urbana-Champaign*

Ensemble-Level Power Management for Dense Blade Servers	66
<i>Parthasarathy Ranganathan, Phil Leech, David Irwin, and Jeffrey Chase</i>	
Techniques for Multicore Thermal Management: Classification and New Exploration	78
<i>James Donald and Margaret Martonosi</i>	
SODA: A Low-Power Architecture for Software Radio	89
<i>Yuan Lin, Hyunseok Leek, Mark Woh, Yoav Harel, Scott Mahlke, Trevor Mudge, Chaitali Chakrabarti, and Krisztián Flautner</i>	

Session 4: Multicore

Chair: Bill Dally, *Stanford University*

An Integrated Framework for Dependable and Revivable Architectures Using Multicore Processors	102
<i>Weidong Shi, Hsien-Hsin S. Lee, Laura Falk, and Mrinmoy Ghosh</i>	
Multiple Instruction Stream Processor	114
<i>Richard A. Hankins, Gautham N. Chinya, Jamison D. Collins, Perry H. Wang, Ryan Rakvic, Hong Wang, and John P. Shen</i>	

Panel Session:

System 2020: Computer Architecture Grand Research Challenges

Mary Jane Irwin, *Penn State University*

David Patterson, *University of California at Berkeley*

John P. Shen, *Intel Corp.*

Keynote 2

Chair: Per Stenström, *Chalmers University of Technology*

The End of Scaling? Revolutions in Technology and Microarchitecture as We Pass the 90 Nanometer Node	128
<i>Philip Emma, IBM T. J. Watson Research Center</i>	

Session 5A: Memory Access Issues

Chair: Antonio González, *Universitat Politècnica de Catalunya/Intel Labs - Barcelona*

Design and Management of 3D Chip Multiprocessors Using Network-in-Memory	130
<i>Feihui Li, Chrysostomos Nicopoulos, Thomas Richardson, Yuan Xie, Vijaykrishnan Narayanan, and Mahmut Kandemir</i>	
Slackened Memory Dependence Enforcement: Combining Opportunistic Forwarding with Decoupled Verification.....	142
<i>Alok Garg, M. Wasiur Rashid, and Michael Huang</i>	

Session 5B: Cache Design I

Chair: Ronny Ronen, *Intel Labs - Haifa*

Balanced-Cache: Reducing Conflict Misses of Direct-Mapped Caches through Programmable Decoders.....	155
<i>Chuanjun Zhang</i>	
A Case for MLP-Aware Cache Replacement	167
<i>Moinuddin K. Qureshi, Daniel N. Lynch, Onur Mutlu, and Yale N. Patt</i>	

Session 6A: Security and Network Processors

Chair: Hsien-Hsin Lee, *Georgia Institute of Technology*

Improving Cost, Performance, and Security of Memory Encryption and Authentication	179
<i>Chenyu Yan, Brian Rogers, Daniel Engleder, Yan Solihin, and Milos Prvulovic</i>	
A Scalable Architecture for High-Throughput Regular-Expression Pattern Matching	191
<i>Benjamin C. Brodie, Ron K. Cytron, and David E. Taylor</i>	
Chisel: A Storage-Efficient, Collision-Free Hash-Based Network Processing Architecture.....	203
<i>Jahangir Hasan, Srihari Cadambi, Venkatta Jakkula, and Srimat Chakradhar</i>	

Session 6B: Multithreading

Chair: Susan Eggers, *University of Washington*

Tolerating Dependences between Large Speculative Threads via Sub-Threads	216
<i>Christopher B. Colohan, Anastassia Ailamaki, J. Gregory Steffan, and Todd C. Mowry</i>	

Bulk Disambiguation of Speculative Threads in Multiprocessors.....227
Luis Ceze, James Tuck, Călin Caşcaval, and Josep Torrellas

Learning-Based SMT Processor Resource Distribution via Hill-Climbing.....239
Seungryul Choi and Donald Yeung

Session 7A: Cache Design II

Chair: Kei Hiraki, *University of Tokyo*

Spatial Memory Streaming.....252
*Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi,
and Andreas Moshovos*

Cooperative Caching for Chip Multiprocessors264
Jichuan Chang and Gurindar S. Sohi

Session 7B: Potpourri

Chair: Daniel Jiménez, *Rutgers University*

Reducing Startup Time in Co-designed Virtual Machines277
Shiliang Hu and James E. Smith

TRAP-Array: A Disk Array Architecture Providing Timely Recovery
to Any Point-in-Time.....289
Qing Yang, Wijun Xiao, and Jin Ren

Session 8A: Dataflow

Chair: Yale Patt, *University of Texas at Austin*

Program Demultiplexing: Data-Flow Based Speculative Parallelization of Methods
In Sequential Programs.....302
Saisanthosh Balakrishnan and Gurindar S. Sohi

Area-Performance Trade-Offs in Tiled Dataflow Architectures314
*Steven Swanson, Andrew Putnam, Martha Mercaldi, Ken Michelson,
Andrew Petersen, Andrew Schwerin, Mark Oskin, and Susan J. Eggers*

Session 8B: Cache Coherence

Chair: Alex Veidenbaum, *University of California, Irvine*

Flexible Snooping: Adaptive Forwarding and Filtering of Snoops in Embedded-Ring
Multiprocessors327
Karin Strauss, Xiaowei Shen, and Josep Torrellas

Interconnect-Aware Coherence Protocols for Chip Multiprocessors339
*Liqun Cheng, Naveen Muralimanohar, Karthik Ramani, Rajeev Balasubramonian,
and John B. Carter*

Keynote 3

Chair: David Kaeli, *Northeastern University*

The Future of Virtualization Technology	352
<i>Steve Herrod, VMware</i>	

Session 9: Quantum Computing

Chair: Matthew Farrens, *University of California, Davis*

Distributed Arithmetic on a Quantum Multicomputer.....	354
<i>Rodney Van Meter, W. J. Munro, Kae Nemoto, and Kohei M. Itoh</i>	
Interconnection Networks for Scalable Quantum Computers	366
<i>Nemanja Isailovic, Yatish Patel, Mark Whitney, and John Kubiawicz</i>	
Quantum Memory Hierarchies: Efficient Designs to Match Available Parallelism in Quantum Computing	378
<i>Darshan D. Thaker, Tzvetan S. Metodi, Andrew W. Cross, Isaac L. Chuang, and Frederic T. Chong</i>	

Author Index	391
---------------------------	-----