

# **2014 ACM/IEEE 41st International Symposium on Computer Architecture**

**(ISCA 2014)**

**Minneapolis, Minnesota, USA  
14 – 18 June 2014**



**IEEE Catalog Number: CFP14030-POD  
ISBN: 978-1-4799-4393-7**

# Table of Contents

## Session 1: Machines and Prototypes

- **Unifying on-chip and inter-node switching within the Anton 2 network.....1**  
Brian Towles, J.P. Grossman, Brian Greskamp, and David E. Shaw (D. E. Shaw Research)
- **A Reconfigurable Fabric for Accelerating Large-Scale Datacenter Services.....13**  
Andrew Putnam (Microsoft), Adrian M. Caulfield (Microsoft), Eric S. Chung (Microsoft), Derek Chiou (Microsoft and University of Texas at Austin), Kypros Constantinides (Amazon), John Demme (Columbia University), Hadi Esmaeilzadeh (Georgia Institute of Technology), Jeremy Fowers (Microsoft), Gopi Prashanth Gopal (Microsoft), Jan Gray (Microsoft), Michael Haselman (Microsoft), Scott Hauck (Microsoft and University of Washington), Stephen Heil (Microsoft), Amir Hormati (Google), Joo-Young Kim (Microsoft), Sitaram Lanka (Microsoft), James Larus (EPFL), Eric Peterson (Microsoft), Simon Pope (Microsoft), Aaron Smith (Microsoft), Jason Thong (Microsoft), Phillip Yi Xiao (Microsoft), Doug Burger (Microsoft)
- **SCORPIO: A 36-Core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-Network Ordering.....25**  
Bhavya K. Daya, Chia-Hsin Owen Chen, Suvinay Subramanian, Woo-Cheol Kwon, Sunghyun Park, Tushar Krishna, Jim Holt, Anantha P. Chandrakasan, Li-Shiuan Peh (Massachusetts Institute of Technology)

## Session 2A: Resilience

- **Avoiding Core's DUE & SDC via Acoustic Wave Detectors and Tailored Error Containment and Recovery.....37**  
Gaurang Upasani (Universitat Politècnica de Catalunya), Xavier Vera (Intel Barcelona Research Center), Antonio González (Universitat Politècnica de Catalunya / Intel Barcelona Research Center)
- **MemGuard: A Low Cost and Energy Efficient Design to Support and Enhance Memory System Reliability.....49**  
Long Chen, Zhao Zhang (Iowa State University)
- **GangES: Gang Error Simulation for Hardware Resiliency Evaluation.....61**  
Siva Kumar Sastry Hari (NVIDIA), Radha Venkatagiri (University of Illinois at Urbana-Champaign), Sarita V. Adve (University of Illinois at Urbana-Champaign), Helia Naeimi (Intel Labs)
- **Real-World Design and Evaluation of Compiler-Managed GPU Redundant Multithreading.....73**  
Jack Wadden (University of Virginia), Alexander Lyashevsky (AMD Research), Sudhanva Gurumurthi (AMD Research), Vilas Sridharan (RAS Architecture, AMD), Kevin Skadron (University of Virginia)

## Session 2B: Design Space Exploration

- **ArchRanker: A Ranking Approach to Design Space Exploration.....85**  
Tianshi Chen (Chinese Academy of Sciences), Qi Guo (Carnegie Mellon University),  
Ke Tang (University of Science and Technology of China), Olivier Temam (Inria),  
Zhiwei Xu (Chinese Academy of Sciences), Zhi-Hua Zhou (Nanjing University),  
Yunji Chen (Chinese Academy of Sciences)
- **Aladdin: A Pre-RTL, Power-Performance Accelerator Simulator  
Enabling Large Design Space Exploration of Customized Architectures.....97**  
Yakun Sophia Shao, Brandon Reagen, Gu-Yeon Wei, David Brooks (Harvard University)
- **SynFull: Synthetic Traffic Models Capturing Cache Coherent Behaviour.....109**  
Mario Badr, Natalie Enright Jerger (University of Toronto)
- **Harnessing ISA Diversity: Design of a Heterogeneous-ISA Chip Multiprocessor.....121**  
Ashish Venkat, Dean M. Tullsen (University of California, San Diego)

## Session 3A: Caches

- **The Direct-to-Data (D2D) Cache: Navigating the Cache Hierarchy with a Single Lookup.....133**  
Andreas Sembrant, Erik Hagersten, David Black-Schaffer (Uppsala University)
- **SC<sup>2</sup>: A Statistical Compression Cache Scheme.....145**  
Angelos Arelakis, Per Stenstrom (Chalmers University of Technology)
- **The Dirty-Block Index.....157**  
Vivek Seshadri (Carnegie Mellon University), Abhishek Bhowmick (Carnegie Mellon University),  
Onur Mutlu (Carnegie Mellon University), Phillip B. Gibbons (Intel Pittsburgh),  
Michael A. Kozuch (Intel Pittsburgh), Todd C. Mowry (Carnegie Mellon University)
- **Going Vertical in Memory Management: Handling Multiplicity by Multi-policy.....169**  
Lei Liu (Chinese Academy of Sciences), Yong Li (University of Pittsburgh),  
Zehan Cui (Chinese Academy of Sciences), Yungang Bao (Chinese Academy of Sciences),  
Mingyu Chen (Chinese Academy of Sciences), Chengyong Wu (Chinese Academy of Sciences)

## Session 3B: GPUs and Parallelism

- **Fine-grain Task Aggregation and Coordination on GPUs.....181**  
Marc S. Orr (University of Wisconsin-Madison / AMD Research), Bradford M. Beckmann (AMD Research),  
Steven K. Reinhardt (AMD Research), David A. Wood (University of Wisconsin-Madison / AMD Research)
- **Enabling Preemptive Multiprogramming on GPUs.....193**  
Ivan Tanasic (Barcelona Supercomputing Center / Universitat Politecnica de Catalunya),  
Isaac Gelado (NVIDIA Research), Javier Cabezas (Barcelona Supercomputing Center / Universitat Politecnica  
de Catalunya), Alex Ramirez (Barcelona Supercomputing Center / Universitat Politecnica de Catalunya),  
Nacho Navarro (Barcelona Supercomputing Center / Universitat Politecnica de Catalunya),  
Mateo Valero (Barcelona Supercomputing Center / Universitat Politecnica de Catalunya)

- **Single-Graph Multiple Flows: Energy Efficient Design Alternative for GPGPUs.....205**  
Dani Voitsechov, Yoav Etsion (Technion – Israel Institute of Technology)
- **HELIX-RC: An Architecture-Compiler Co-Design for Automatic Parallelization of Irregular Programs.....217**  
Simone Campanoni (Harvard University), Kevin Brownell (Harvard University), Svilen Kanev (Harvard University), Timothy M. Jones (University of Cambridge), Gu-Yeon Wei (Harvard University), David Brooks (Harvard University)

## Session 4: Emerging Technologies

- **Efficient Digital Neurons for Large Scale Cortical Architectures.....229**  
James E. Smith (University of Wisconsin-Madison)
- **An Examination of the Architecture and System-level Tradeoffs of Employing Steep Slope Devices in 3D CMPs.....241**  
Karthik Swaminathan, Huichu Liu, Jack Sampson, Vijaykrishnan Narayanan (Pennsylvania State University)
- **STAG: Spintronic-Tape Architecture for GPGPU Cache Hierarchies.....253**  
Rangharajan Venkatesan, Shankar Ganesh Ramasubramanian, Swagath Venkataramani, Kaushik Roy, Anand Raghunathan (Purdue University)

## Session 5A: NVRAM

- **Memory Persistency.....265**  
Steven Pelley, Peter M. Chen, Thomas F. Wenisch (University of Michigan)
- **Reducing Access Latency of MLC PCMs through Line Striping.....277**  
Morteza Hoseinzadeh (Sharif University of Technology), Mohammad Arjomand (Sharif University of Technology), Hamid Sarbazi-Azad (Sharif University of Technology / Institute for Research in Fundamental Sciences)
- **HIOS: A Host Interface I/O Scheduler for Solid State Disks.....289**  
Myoungsoo Jung (University of Texas at Dallas), Wonil Choi (University of Texas at Dallas), Shekhar Srikantiah (Qualcomm), Joonhyuk Yoo (Daegu University), Mahmut T. Kandemir (Pennsylvania State University)

## Session 5B: Datacenters and Cloud

- **Towards Energy Proportionality for Large-Scale Latency-Critical Workloads.....301**  
David Lo (Stanford University), Liqun Cheng (Google), Rama Govindaraju (Google), Luiz André Barroso (Google), Christos Kozyrakis (Stanford University)
- **SleepScale: Runtime Joint Speed Scaling and Sleep States Management for Power Efficient Data Centers.....313**  
Yanpei Liu (University of Wisconsin-Madison), Stark C. Draper (University of Toronto), Nam Sung Kim (University of Wisconsin-Madison)

- **Optimizing Virtual Machine Consolidation Performance on NUMA Server Architecture for Cloud Workloads.....325**  
Ming Liu, Tao Li (University of Florida)

## Session 6A: DRAM

- **Row-Buffer Decoupling: A Case for Low-Latency DRAM Microarchitecture.....337**  
Seongil O (Seoul National University), Young Hoon Son (Seoul National University),  
Nam Sung Kim (University of Wisconsin-Madison), Jung Ho Ahn (Seoul National University)
- **Half-DRAM: a High-bandwidth and Low-power DRAM Architecture from the Rethinking of Fine-grained Activation.....349**  
Tao Zhang (Pennsylvania State University / NVIDIA), Ke Chen (Oracle),  
Cong Xu (Pennsylvania State University), Guangyu Sun (Peking University), Tao Wang (Peking University),  
Yuan Xie (Pennsylvania State University)
- **Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors.....361**  
Yoongu Kim (Carnegie Mellon University), Ross Daly, Jeremie Kim (Carnegie Mellon University),  
Chris Fallin, Ji Hye Lee (Carnegie Mellon University), Donghyuk Lee (Carnegie Mellon University),  
Chris Wilkerson (Intel Labs), Konrad Lai, Onur Mutlu (Carnegie Mellon University)

## Session 6B: Circuits and Architecture

- **Architecture Implications of Pads as a Scarce Resource.....373**  
Runjie Zhang (University of Virginia), Ke Wang (University of Virginia), Brett H. Meyer (McGill University),  
Mircea R. Stan (University of Virginia), Kevin Skadron (University of Virginia)
- **Increasing Off-Chip Bandwidth in Multi-Core Processors with Switchable Pins.....385**  
Shaoming Chen, Yue Hu, Ying Zhang, Lu Peng, Jesse Ardonne, Samuel Irving,  
Ashok Srivastava (Louisiana State University)
- **A Low Power and Reliable Charge Pump Design for Phase Change Memories.....397**  
Lei Jiang, Bo Zhao, Jun Yang, Youtao Zhang (University of Pittsburgh)

## Session 7A: Coherence and Replay

- **Fractal++: Closing the Performance Gap between Fractal and Conventional Coherence.....409**  
Gwendolyn Voskuilen, T. N. Vijaykumar (Purdue University)
- **OmniOrder: Directory-Based Conflict Serialization of Transactions.....421**  
Xuehai Qian (University of California, Berkeley), Benjamin Sahelices (Universidad de Valladolid),  
Josep Torrellas (University of Illinois at Urbana-Champaign)
- **Pacifier: Record and Replay for Relaxed-Consistency Multiprocessors with Distributed Directory Protocol.....433**  
Xuehai Qian (University of California, Berkeley), Benjamin Sahelices (Universidad de Valladolid),  
Depei Qian (Beihang University)

- **Replay Debugging: Leveraging Record and Replay for Program Debugging**.....445  
Nima Honarmand, Josep Torrellas (University of Illinois at Urbana-Champaign)

## Session 7B: Security/OOO Processors

- **The CHERI capability model: Revisiting RISC in an age of risk**.....457  
Jonathan Woodruff (University of Cambridge), Robert N. M. Watson (University of Cambridge),  
David Chisnall (University of Cambridge), Simon W. Moore (University of Cambridge),  
Jonathan Anderson (University of Cambridge), Brooks Davis (SRI International), Ben Laurie (Google UK Ltd),  
Peter G. Neumann (SRI International), Robert Norton (University of Cambridge),  
Michael Roe (University of Cambridge)
- **CODOMs: Protecting Software with Code-centric Memory Domains**.....469  
Lluís Vilanova (Barcelona Supercomputing Center / Universitat Politècnica de Catalunya /  
Technion – Israel Institute of Technology), Muli Ben-Yehuda (Technion - Israel Institute of Technology),  
Nacho Navarro (Barcelona Supercomputing Center / Universitat Politècnica de Catalunya),  
Yoav Etsion (Technion – Israel Institute of Technology),  
Mateo Valero (Barcelona Supercomputing Center / Universitat Politècnica de Catalunya)
- **EOLE: Paving the Way for an Effective Implementation of Value Prediction**.....481  
Arthur Perais, André Seznec (IRISA/INRIA)
- **Improving the Energy Efficiency of Big Cores**.....493  
Kenneth Czechowski (Georgia Institute of Technology), Victor W. Lee (Intel), Ed Grochowski (Intel),  
Ronny Ronen (Intel), Ronak Singhal (Intel), Richard Vuduc (Georgia Institute of Technology),  
Pradeep Dubey (Intel)

## Session 8: Accelerators

- **General-Purpose Code Acceleration with Limited-Precision Analog Computation**.....505  
Renée St. Amant (University of Texas at Austin), Amir Yazdanbakhsh (Georgia Institute of Technology),  
Jongse Park (Georgia Institute of Technology), Bradley Thwaites (Georgia Institute of Technology),  
Hadi Esmaeilzadeh (Georgia Institute of Technology), Arjang Hassibi (University of Texas at Austin),  
Luis Ceze (University of Washington), Doug Burger (Microsoft Research)
- **Race Logic: A Hardware Acceleration for Dynamic Programming Algorithms**.....517  
Advait Madhavan, Timothy Sherwood, Dmitri Strukov (University of California, Santa Barbara)
- **Eliminating Redundant Fragment Shader Executions on a Mobile GPU via Hardware Memoization**...529  
Jose-Maria Arnau (Universitat Politècnica de Catalunya),  
Joan-Manuel Parcerisa (Universitat Politècnica de Catalunya), Polychronis Xekalakis (Intel)
- **WebCore: Architectural Support for Mobile Web Browsing**.....541  
Yuhao Zhu, Vijay Janapa Reddi (The University of Texas at Austin)