Sunday, December 2, 2012

18:00-22:00: Conference Registration and Welcome Reception [Roof]

Monday, December 3, 2012

7:30-18:30: Conference Registration [British Columbia Foyer]

7:30-8:30: Continental Breakfast [British Columbia Foyer]

8:30-9:00: Welcome and Opening Remarks by General and Program Chairs [Columbia Ballroom]

9:00-10:00: Keynote I [Columbia Ballroom]

Charles Webb, IBM
The Next Era of Enterprise Processor Design

10:00-10:30: Break

10:30-12:00: Lightning Session [Columbia Ballroom]
Session Chair: Onur Mutlu (Carnegie Mellon)

100-second "key idea" pitch for each paper

12:00-13:30: Lunch [Pacific Ballroom]

13:30-15:00: Session IA - Memory Systems I [Columbia Ballroom]
Session Chair: Eiman Ebrahimi (NVIDIA)

FPB: Fine-grained Power Budgeting to Improve Write Throughput of Multi-level Cell Phase Change Memory
Lei Jiang, Youtao Zhang, Bruce R. Childers, and Jun Yang (University of Pittsburgh)

Leveraging Heterogeneity in DRAM Main Memories to Accelerate Critical Word Access
Niladrish Chatterjee, Manjunath Shevgoor, Rajeev Balasubramonian, and Al Davis (University of Utah), Zhen Fang (Nvidia), Ramesh Illikkal, and Ravi Iyer (Intel)

Transactional Memory Architecture and Implementation for IBM System z
Christian Jacobi, Timothy Slegel, and Dan Greiner (IBM)

13:30-15:00: Session IB - Fault Tolerance [British Ballroom]
Session Chair: Ahmed Louri (University of Arizona and NSF)

Warped-DMR: Light-weight Error Detection for GPGPU
Hyeran Jeon and Murali Annavaram (University of Southern California)

The Performance Vulnerability of Architectural and Non-architectural Arrays to Permanent Faults
Damien Hardy (University of Cyprus and University of Rennes 1, IRISA), Isidoros

Sideris, Nikolas Ladas, and Yiannakis Sazeides (University of Cyprus)

NoCAlert: An On-Line and Real-Time Fault Detection Mechanism for Network-on-Chip Architectures
Andreas Prodromou, Andreas Panteli, Chrysostomos Nicopoulos, and Yiannakis Sazeides (University of Cyprus)

15:00-15:30: Break

15:30-17:30: Session IIA - GPUs and SIMD [Columbia Ballroom]
Session Chair: Hyesoon Kim (Georgia Tech)

Cache-Conscious Wavefront Scheduling
Timothy G. Rogers (University of British Columbia), Mike O'Connor (AMD Research), and Tor M. Aamodt (University of British Columbia)

Libra: Tailoring SIMD Execution using Heterogeneous Hardware and Dynamic Configurability
Yongjun Park, Jason Jong Kyu Park, Hyunchul Park, and Scott Mahlke (University of Michigan, Ann Arbor)

Unifying Primary Cache, Scratch, and Register File Memories in a Throughput Processor
Mark Gebhart (The University of Texas at Austin), Stephen W. Keckler (NVIDIA / The University of Texas at Austin), Brucek Khailany and Ronny Krashinsky (NVIDIA), and William J. Dally (NVIDIA and Stanford University)

Kernel Weaver: Automatically Fusing Database Primitives for Efficient GPU Computation
Haicheng Wu (Georgia Institute of Technology), Gregory Diamos (NVIDIA Research), Srihari Cadambi (NEC Laboratories America), and Sudhakar Yalamanchili (Georgia Institute of Technology)

15:30-17:30: Session IIB - Energy I [Brisith Ballroom]
Session Chair: Boris Grot (EPFL)

KnightShift: Scaling the Energy Proportionality Wall Through Server-level Heterogeneity
Daniel Wong and Murali Annavaram (University of Southern California)

Rethinking DRAM Powermodes for Energy Proportionality
Krishna T. Malladi (Stanford University), Ian Shaeffer and Liji Gopalakrishnan (Rambus Inc), David Lo (Stanford University), Benjamin C. Lee (Duke University), and Mark Horowitz (Stanford University)

CoScale: Coordinating CPU and Memory System DVFS in Server Systems
Qingyuan Deng (Rutgers University), David Meisner (Facebook Inc.), Abhishek Bhattacharjee (Rutgers University), Thomas F. Wenisch (University of Michigan), and Ricardo Bianchini (Rutgers University)

Predicting Performance Impact of DVFS for Realistic Memory Systems
Rustam Miftakhutdinov (The University of Texas at Austin), Eiman Ebrahimi (NVIDIA), and Yale N. Patt (The University of Texas at Austin)

17:30-17:45: Break

17:45-18:15: Program Chair's Remarks [Columbia Ballroom]

Paper Selection Process for Micro-45

19:00-24:00: Conference Excursion and Banquet - Vancouver Aquarium

Buses depart to the Aquarium from the Breezeway approximately every 10 minutes from 18:45 to 19:45; buses return to the hotel approximately every 20 minutes between 22:00 and 24:00.

Tuesday, December 4, 2012

7:30-18:30: Conference Registration [British Columbia Foyer]

7:30-8:30: Continental Breakfast [British Columbia Foyer]

8:30-10:00: Session IIIA - Big Data [Columbia Ballroom]
Session Chair: Yoav Etsion (Technion)

Vector Extensions for Decision Support DBMS Acceleration
Timothy Hayes, Oscar Palomar, Osman Unsal, Adrian Cristal, and Mateo Valero (Barcelona Supercomputing Center)

NOC-Out: Microarchitecting a Scale-Out Processor
Pejman Lotfi-Kamran, Boris Grot, and Babak Falsafi (EPFL)

SLICC: Self-Assembly of Instruction Cache Collectives for OLTP Workloads
Islam Atta (University of Toronto), Pinar Tözün and Anastasia Ailamaki (EPFL), and Andreas Moshovos (University of Toronto)

8:30-10:00: Session IIIB - Energy II [British Ballroom]
Session Chair: Niti Madan (Oracle)

Systematic Energy Characterization of CMP/SMT Processor Systems via Automated Micro-Benchmarks
Ramon Bertran (Barcelona Supercomputing Center and IBM Research), Alper Buyuktosunoglu and Meeta Sharma Gupta (IBM Research), Marc Gonzalez (Barcelona Supercomputing Center), and Pradip Bose (IBM Research)

AUDIT: Stress Testing the Automatic Way
Youngtaek Kim and Lizy Kurian John (The University of Texas at Austin), Sanjay Pant (AMD), Srilatha Manne and Michael Schulte (AMD Research), and W. Lloyd Bircher and Madhu S. Sibi Govindan (AMD)

Accurate Fine-Grained Processor Power Proxies
Wei Huang (AMD), Charles Lefurgy (IBM Research), William Kuk (Purdue University), Alper Buyuktosunoglu (IBM Research), Michael Floyd (IBM System and Technology Group), Karthick Rajamani and Malcolm Allen-Ware (IBM Research), and Bishop Brock (IBM System and Technology Group)

10:00-10:30: Break

10:30-12:00: Session IVA - Memory Systems II [Columbia Ballroom]
Session Chair: Andreas Moshovos (Toronto)

Fundamental Latency Trade-offs in Architecting DRAM Caches

Moinuddin Qureshi (Georgia Institute of Technology), Gabriel H. Loh (AMD Research)

A Mostly-Clean DRAM Cache for Effective Hit Speculation and Self-Balancing Dispatch

Jaewoong Sim (Georgia Institute of Technology), Gabriel H. Loh (AMD Research), Hyesoon Kim (Georgia Institute of Technology), Mike O'Connor (AMD Research), and Mithuna Thottethodi (AMD Research / Purdue University)

CoLT: Coalesced Large-Reach TLBs
Binh Pham and Viswanathan Vaidyanathan (Rutgers University), Aamer Jaleel (Intel VSSAD), and Abhishek Bhattacharjee (Rutgers University)

10:30-12:00: Session IVB - Interconnects [British Ballroom]
Session Chair: Avinash Kodi (Ohio)

NoRD: Node-Router Decoupling for Effective Power-gating of On-Chip Routers
Lizhong Chen and Timothy Mark Pinkston (University of Southern California)

Dynamic Reconfiguration of 3D Photonic On-chip Interconnects for Maximizing Performance and Improving Fault Tolerance
Randy Morris and Avinash Kodi (Ohio University) and Ahmed Louri (University of Arizona/NSF)

Addressing End-to-End Memory Access Latency in NoC-Based Multicores
Akbar Sharifi, Emre Kultursay, Mahmut Kandemir, and Chita R. Das (The Pennsylvania State University)

12:00-13:30: Lunch and Bob Rau Award Presentation [Pacific Ballroom]

12:30-13:15: Bob Rau Award Presentation - Speech by Josh Fisher

13:30-15:00: Session VA - Core Design [Columbia Ballroom]
Session Chair: Brian Prasky (IBM)

MorphCore: An Energy-Efficient Microarchitecture for High Performance ILP and High Throughput TLP
Khubaib (UT Austin), Aater Suleman (Calxeda/HPS), Milad Hashemi (UT Austin), Chris Wilkerson (Intel Labs), and Yale Patt (UT Austin)

Composite Cores: Pushing Heterogeneity into a Core
Andrew Lukefahr, Shruti Padmanabha, Reetuparna Das, Faissal M. Sleiman, Ronald Dreslinski, Thomas F. Wenisch, and Scott Mahlke (University of Michigan)

Control-Flow Decoupling
Rami Sheikh, James Tuck, and Eric Rotenberg (North Carolina State University)

13:30-15:00: Session VB - Coherence and Consistency [British Ballrooom]
Session Chair: Satish Narayanasamy (Michigan)

Spatiotemporal Coherence Tracking
Mohammad Alisafaee (EPFL)

Predicting Coherence Communication by Tracking Synchronization Points at Run Time
Socrates Demetriades and Sangyeun Cho (University of Pittsburgh)

Vulcan: Hardware Support for Detecting Sequential Consistency Violations Dynamically
Abdullah Muzahid, Shanxiang Qi, and Josep Torrellas (University of Illinois at Urbana-Champaign)

15:00-17:30: Poster Session and Reception [British Ballroom, British Columbia Foyer]
Session Chair: Onur Mutlu (Carnegie Mellon)

An open forum with all conference papers presented on posters with authors available for discussion. Drinks and snacks served.

17:30-19:30: Micro Business Meeting [Columbia Ballroom]

17:30 - 18:30 Drinks served [British Columbia Foyer]

Wednesday, December 5, 2012

7:30-12:30: Conference Registration [British Columbia Foyer]

7:30-8:30: Continental Breakfast [British Columbia Foyer]

8:30-9:30: Keynote II [Columbia Ballroom]

Turner Whitted, Microsoft Research
Architectures for Smart Displays

9:30-10:00: Break

10:00-11:00: Session VIA - Caching [Columbia Ballroom]
Session Chair: Gabriel Loh (AMD Research)

Amoeba-Cache: Adaptive Blocks for Eliminating Waste in the Memory Hierarchy
Snehasish Kumar (Simon Fraser University), Hongzhou Zhao (University of Rochester), Arrvindh Shriraman and Eric Matthews (Simon Fraser University), Sandhya Dwarkadas (University of Rochester), and Lesley Shannon (Simon Fraser University)

Improving Cache Management Policies Using Dynamic Reuse Distances
Nam Duong, Dali Zhao, Taesu Kim, Rosario Cammarota, and Alexander Veidenbaum (UC Irvine) and Mateo Valero (UPC)

10:00-11:00: Session VIB - Modeling and Partitioning [British Ballroom]
Session Chair: Dam Sunwoo (ARM)

Kernel Partitioning of Streaming Applications: A Statistical Approach to an NP-complete Problem
Petar Radojkovic and Paul M. Carpenter (Barcelona Supercomputing Center), Miquel Moreto (Universitat Politecnica de Catalunya and Barcelona Supercomputing Center), Alex Ramirez (Barcelona Supercomputing Center), and Francisco J. Cazorla (Barcelona Supercomputing Center and Spanish National Research Council)

Inferred Models for Dynamic and Sparse Hardware-Software Spaces
Weidan Wu and Benjamin C. Lee (Duke University)

11:00-11:30: Break

11:30-12:30: Session VIIA - Dynamic Optimization and Parallelization [Columbia Ballroom]
Session Chair: David Hansquine (Qualcomm)

SMARTQ: Software-Managed Alias Register Queue for Dynamic Optimizations
Cheng Wang, Youfeng Wu, Hongbo Rong, and Hyunchul Park (Intel Labs)

Profiling Data-Dependence to Assist Parallelization: Framework, Scope, and Optimization
Alain Ketterlin and Philippe Clauss (INRIA & Université de Strasbourg)

11:30-12:30: Session VIIB - Accelerators [British Ballroom]
Session Chair: Reetuparna Das (Michigan)

Neural Acceleration for General-Purpose Approximate Programs
Hadi Esmaeilzadeh, Adrian Sampson, and Luis Ceze (University of Washington) and Doug Burger (Microsoft Research)

Designing a Programmable Wire-Speed Regular-Expression Matching Accelerator
Jan van Lunteren, Christoph Hagleitner, and Kubilay Atasu (IBM Research), Giora Biran, Uzi Shvadron, and Timothy Heil (IBM)

12:30-12:45: Closing Remarks [Columbia Ballroom]