

2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2015)

**Waikiki, Hawaii, USA
5 – 9 December 2015**



**IEEE Catalog Number: CFP15071-POD
ISBN: 978-1-5090-6601-8**

**Copyright © 2015, The Association for Computing Machinery, Inc. (ACM)
All Rights Reserved**

****** This is a print representation of what appears in the IEEE Digital Library. Some format issues inherent in the e-media version may also appear in this print version.***

IEEE Catalog Number:	CFP15071-POD
ISBN (Print-On-Demand):	978-1-5090-6601-8
ISBN (Online):	978-1-4503-4034-2
ISSN:	1072-4451

Additional Copies of This Publication Are Available From:

Curran Associates, Inc
57 Morehouse Lane
Red Hook, NY 12571 USA
Phone: (845) 758-0400
Fax: (845) 758-2633
E-mail: curran@proceedings.com
Web: www.proceedings.com

CURRAN ASSOCIATES INC.
proceedings
.com

Table of Contents

MICRO-48 Organization	xiii
MICRO-48 Program Committee	xiv
MICRO-48 External Review Committee	xvi

Best Paper Candidates

- **Large Pages and Lightweight Memory Management in Virtualized Environments: Can You Have it Both Ways?** **1**
Binh Pham (Rutgers University), Jan Vesely (Rutgers University), Gabriel H. Loh (AMD Research), Abhishek Bhattacharjee (Rutgers University)
- **Exploiting Commutativity to Reduce the Cost of Updates to Shared Data in Cache-Coherent Systems.....** **13**
Guowei Zhang (Massachusetts Institute of Technology), Webb Horn (Massachusetts Institute of Technology), Daniel Sanchez (Massachusetts Institute of Technology)
- **CCICheck: Using μ hb Graphs to Verify the Coherence-Consistency Interface** **26**
Yatin A. Manerkar (Princeton University), Daniel Lustig (Princeton University), Michael Pellauer (NVIDIA), Margaret Martonosi (Princeton University)

Session A1: Cache

- **HyComp: A Hybrid Cache Compression Method for Selection of Data-Type-Specific Compression Methods** **38**
Angelos Arelakis (Chalmers University of Technology), Fredrik Dahlgren (Chalmers University of Technology), Per Stenstrom (Chalmers University of Technology)
- **Doppelganger: A Cache for Approximate Computing** **50**
Joshua San Miguel, (University of Toronto), Jorge Albericio (University of Toronto), Andreas Moshovos (University of Toronto), Natalie Enright Jerger (University of Toronto)
- **The Application Slowdown Model: Quantifying and Controlling the Impact of Inter-Application Interference at Shared Caches and Main Memory.....** **62**
Lavanya Subramanian (Intel Labs), Vivek Seshadri (Carnegie Mellon University), Arnab Ghosh (IIT Kanpur), Samira Khan (University of Virginia), Onur Mutlu (Carnegie Mellon University)
- **MORC: A Manycore-Oriented Compressed Cache** **76**
Tri M. Nguyen (Princeton University), David Wentzlaff (Princeton University)

Session A2: Security

- **Avoiding Information Leakage in the Memory Controller with Fixed Service Policies** **89**
Ali Shafiq (University of Utah), Akhila Gundu (University of Utah), Manjunath Shevgoor (University of Utah), Rajeev Balasubramonian (University of Utah), Mohit Tiwari (University of Texas at Austin)
- **Fork Path: Improving Efficiency of ORAM by Removing Redundant Memory Accesses** **102**
Xian Zhang (Peking University), Guangyu Sun (Peking University), Chao Zhang (Peking University), Weiqi Zhang (Peking University), Yun Liang (Peking University), Tao Wang (Peking University), Yiran Chen (University of Pittsburgh), Jia Di (University of Arkansas)
- **Locking Down Insecure Indirection with Hardware-Based Control-Data Isolation** **115**
William Arthur (University of Michigan), Sahil Madeka (University of Michigan), Reetuparna Das (University of Michigan), Todd Austin (University of Michigan)
- **Authenticache: Harnessing Cache ECC for System Authentication** **128**
Anys Bacha (The Ohio State University), Radu Teodorescu (The Ohio State University)

Session B1: Prefetching

- **Efficiently Prefetching Complex Address Patterns** 141
Manjunath Shevgoor (University of Utah), Sahil Koladiya (University of Utah), Rajeep Balasubramonian (University of Utah), Chris Wilkerson (Intel Labs), Seth Pugsley (Intel Labs), Zeshan Chishti (Intel Labs)
- **Self-Contained, Accurate Precomputation Prefetching** 153
Islam Atta (University of Toronto), Xin Tong (University of Toronto), Vijayalakshmi Srinivasan (IBM), Ioana Baldini (IBM), Andreas Moshovos (University of Toronto)
- **Confluence: Unified Instruction Supply for Scale-Out Servers** 166
Cansu Kaynak (EcoCloud, EPFL), Boris Grot (University of Edinburgh), Babak Falsafi (EcoCloud, EPFL)
- **IMP: Indirect Memory Prefetcher** 178
Xiangyao Yu (MIT), Christopher J. Hughes (Intel), Nadathur Satish (Intel), Srinivas Devadas (Massachusetts Institute of Technology)

Session B2: Concurrency

- **DeSC: Decoupled Supply-Compute Communication Management for Heterogeneous Architectures**..... 191
Tae Jun Ham (Princeton University), Juan Luis Aragon (University of Murcia), Margaret Martonosi (Princeton University)
- **Efficient Warp Execution in Presence of Divergence with Collaborative Context Collection** 204
Farzad Khorasani (University of California Riverside), Rajiv Gupta (University of California Riverside), Laxmi N. Bhuyan (University of California Riverside)
- **Control Flow Coalescing on a Hybrid Dataflow/von Neumann GPGPU** 216
Dani Voitsechov (Technion), Yoav Etsion (Technion)
- **A Scalable Architecture for Ordered Parallelism** 228
Mark C. Jeffrey (Massachusetts Institute of Technology), Suvinay Subramanian (Massachusetts Institute of Technology), Cong Yan (Massachusetts Institute of Technology), Joel Emer (Massachusetts Institute of Technology / NVIDIA), Daniel Sanchez (Massachusetts Institute of Technology)

Session C1: DRAM

- **More is Less: Improving the Energy Efficiency of Data Movement via Opportunistic Use of Sparse Codes** 242
Yanwei Song (University of Rochester), Engin Ipek (University of Rochester)
- **Improving DRAM Latency with Dynamic Asymmetric Subarray** 255
Shih-Lien Lu (Intel), Ying-Chen Lin (National Taiwan University), Chia-Lin Yang (National Taiwan University)
- **Gather-Scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-unit Strided Accesses** 267
Vivek Seshadri (Carnegie Mellon University), Thomas Mullins (Carnegie Mellon University), Amirali Boroumand (Carnegie Mellon University), Onur Mutlu (Carnegie Mellon University), Phillip B. Gibbons (Carnegie Mellon University), Michael A. Kozuch (Intel Pittsburgh), Todd C. Mowry (Carnegie Mellon University)

Session C2: Voltage

- **The CRISP Performance Model for Dynamic Voltage and Frequency Scaling in a GPGPU** 281
Rajib Nath (UCSD), Dean Tullsen (UCSD)
- **Safe Limits on Voltage Reduction Efficiency in GPUs: a Direct Measurement Approach** 294
Jingwen Leng (IBM T.J. Watson Research), Alper Buyuktosunoglu (IBM T.J. Watson Research), Ramon Bertran (IBM T.J. Watson Research), Pradip Bose (IBM T.J. Watson Research), Vijay Janapa Reddi (University of Texas at Austin)

- **Adaptive Guardband Scheduling to Improve System-level Efficiency of POWER7+** 308
Yazhou Zu (University of Texas at Austin), Charles R. Lefurgy (IBM), Jingwen Leng (University of Texas at Austin), Matthew Halpern (University of Texas at Austin), Michael S. Floyd (IBM), Vijay Janapa Reddi (University of Texas at Austin)

Session D1: Micro-architecture

- **DynaMOS: Dynamic Schedule Migration for Heterogeneous Cores** 322
Shruti Padmanabha (University of Michigan), Andrew Lukefahr (University of Michigan), Reetuparna Das (University of Michigan), Scott Mahlke (University of Michigan)
- **Long Term Parking: Criticality-aware Resource Allocation in OOO Processors** 334
Andreas Sembrant (Uppsala University), Trevor E. Carlson (Uppsala University), Erik Hagersten (Uppsala University), David Black-Schaffer (Uppsala University), Arthur Perais (IRISA/INRIA), Andre Sez nec (IRISA/INRIA), Pierre Michaud (IRISA/INRIA)
- **The Inner Most Loop Iteration counter: a new dimension in branch history** 347
Andre Sez nec (IRISA/INRIA), Joshua San Miguel (University of Toronto), Jorge Albericio (University of Toronto)
- **Filtered Runahead Execution with a Runahead Buffer** 358
Milad Hashemi (University of Texas at Austin), Yale N. Patt (University of Texas at Austin)
- **Bungee Jumps: Accelerating Indirect Branches Through HW/SW Co-Design** 370
Daniel S. McFarlin (Carnegie Mellon University), Craig Zilles (University of Illinois, Urbana-Champaign)

Session D2: GPU

- **SAWS: Synchronization Aware GPGPU Warp Scheduling for Multiple Independent Warp Schedulers** 383
Jiwei Liu (University of Pittsburgh), Jun Yang (University of Pittsburgh), Rami Melhem (University of Pittsburgh)
- **Enabling Coordinated Register Allocation and Thread-level Parallelism Optimization for GPUs**.....395
Xiaolong Xie (Center for Energy-Efficient Computing and Applications, Peking University), Yun Liang (Center for Energy-Efficient Computing and Applications, Peking University), Xiuhong Li (Center for Energy-Efficient Computing and Applications, Peking University), Yudong Wu (Center for Energy-Efficient Computing and Applications, Peking University), Guangyu Sun (Center for Energy-Efficient Computing and Applications, Peking University), Tao Wang (Center for Energy-Efficient Computing and Applications, Peking University), Dongrui Fan (Institute of Computing Technology, Chinese Academy of Sciences)
- **Free Launch: Optimizing GPU Dynamic Kernel Launches through Thread Reuse** 407
Guoyang Chen (North Carolina State University), Xipeng Shen (North Carolina State University)
- **GPU Register File Virtualization** 420
Hyeran Jeon (San Jose State University), Gokul Subramanian Ravi (University of Wisconsin-Madison), Nam Sung Kim (University of Illinois, Urbana-Champaign), Murali Annavaram (University of Southern California)
- **WarpPool: Sharing Requests with Inter-Warp Coalescing for Throughput Processors** 433
John Kloosterman (University of Michigan, Ann Arbor), Jonathan Beaumont (University of Michigan, Ann Arbor), Mick Wollman (University of Michigan, Ann Arbor), Ankit Sethia (University of Michigan, Ann Arbor), Ron Dreslinski (University of Michigan, Ann Arbor), Trevor Mudge (University of Michigan, Ann Arbor), Scott Mahlke (University of Michigan, Ann Arbor)

Session E1: Accelerator

- **Ultra-Low Power Render-Based Collision Detection for CPU/GPU Systems** 445
Enrique de Lucas (Polytechnic University of Catalonia), Pedro Marcuello (Broadcom Corporation), Joan-Manuel Parcerisa (Polytechnic University of Catalonia), Antonio Gonzalez (Polytechnic University of Catalonia)

- **Execution Time Prediction for Energy-Efficient Hardware Accelerators** 457
Tao Chen (Cornell University), Alexander Rucker (Cornell University), G. Edward Suh (Cornell University)
- **Border Control: Sandboxing Accelerators** 470
Lena E. Olson (University of Wisconsin-Madison), Jason Power (University of Wisconsin-Madison), Mark D. Hill (University of Wisconsin-Madison), David A. Wood (University of Wisconsin-Madison)
- **Neural Acceleration for GPU Throughput Processors** 482
Amir Yazdanbakhsh (Georgia Institute of Technology), Jongse Park (Georgia Institute of Technology), Hardik Sharma (Georgia Institute of Technology), Pejman Lotfi-Kamran (IPM - Institute for Research in Fundamental Sciences), Hadi Esmaeilzadeh (Georgia Institute of Technology)
- **Neuromorphic Accelerators: A Comparison Between Neuroscience and Machine-Learning Approaches** 494
Zidong Du (State Key Laboratory of Computer Architecture, ICT, CAS, China), Daniel D Ben-Dayan Rubin (Intel Labs, Israel), Yunji Chen (State Key Laboratory of Computer Architecture, ICT, CAS, China), Liqiang He (College of Computer Science, Inner Mongolia University, China), Tianshi Chen (State Key Laboratory of Computer Architecture, ICT, CAS, China), Lei Zhang (University of Science and Technology of China), Chengyong Wu (State Key Laboratory of Computer Architecture, ICT, CAS, China), Olivier Temam (Inria, France)

Session E2: Mobile & Emerging Systems

- **Prediction-Guided Performance-Energy Trade-off for Interactive Applications** 508
Daniel Lo (Cornell University), Taejoon Song (Cornell University), G. Edward Suh (Cornell University)
- **Architecture-aware Automatic Computation Offload for Native Applications** 521
Gwangmu Lee (POSTECH), Hyunjoon Park (POSTECH), Seonyeong Heo (POSTECH), Kyung-Ah Chang (Samsung Electronics), Hyogun Lee (Samsung Electronics), Hanjun Kim (POSTECH)
- **Fast Support for Unstructured Data Processing: the Unified Automata Processor** 533
Yuanwei Fang (University of Chicago), Tung T. Hoang (University of Chicago), Michela Becchi (University of Missouri), Andrew A. Chien (University of Chicago)
- **Enabling Interposer-based Disintegration of Multi-core Processors** 546
Ajaykumar Kannan (University of Toronto), Natalie Enright Jerger (University of Toronto), Gabriel H. Loh (AMD Research)
- **DCS: A Fast and Scalable Device-Centric Server Architecture** 559
Jaehyung Ahn (POSTECH), Dongup Kwon (POSTECH), Youngsok Kim (POSTECH), Mohammadamin Ajdari (POSTECH), Jaewon Lee (POSTECH), Jangwoo Kim (POSTECH)

Session F1: Datacenter

- **Modeling the Implications of DRAM Failures and Protection Techniques on Datacenter TCO** 572
Panagiota Nikolaou (University of Cyprus), Yiannakis Sazeides (University of Cyprus), Lorena Ndreu (University of Cyprus), Marios Kleanthous (MAP S.Platis)
- **TimeTrader: Exploiting Latency Tail to Save Datacenter Energy for Online Search** 585
Balajee Vamanan (Purdue University), Hamza Bin Sohail (Purdue University), Jahangir Hasan (Google Inc.), T. N. Vijaykumar (Purdue University)
- **Rubik: Fast Analytical Power Management for Latency-Critical Systems** 598
Harshad Kasture (MIT), Davide B. Bartolini (MIT), Nathan Beckmann (MIT), Daniel Sanchez (MIT)

Session F2: Memory Systems

- **CLEAN-ECC: High Reliability ECC for Adaptive Granularity Memory System** 611
Seong-Lyong Gong (University of Texas at Austin), Minsoo Rhu (NVIDIA Research), Jungrae Kim (University of Texas at Austin), Jinsuk Chung (University of Texas at Austin), Mattan Erez (University of Texas at Austin)

- **vCache: Architectural Support for Transparent and Isolated Virtual LLCs in Virtualized Environments** 623
Daehoon Kim (University of Illinois at Urbana-Champaign), Hwanju Kim (University of Cambridge), Nam Sung Kim (University of Illinois at Urbana-Champaign), Jaehyuk Huh (KAIST)
- **An integrated concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors** 635
Kathryn E. Gray (University of Cambridge), Gabriel Kerneis (University of Cambridge), Dominic Mulligan (University of Cambridge), Christopher Pulte (University of Cambridge), Susmit Sarkar (University of St. Andrews), Peter Sewell (University of Cambridge)

Session G1: Coherence, Consistency, Persistency

- **Efficient GPU Synchronization without Scopes: Saying No to Complex Consistency Models** 647
Matthew D. Sinclair (University of Illinois at Urbana-Champaign), Johnathan Alsop (University of Illinois at Urbana-Champaign), Sarita V. Adve (University of Illinois at Urbana-Champaign)
- **Efficient Persist Barriers for Multicores** 660
Arpit Joshi (University of Edinburgh), Vijay Nagarajan (University of Edinburgh), Marcelo Cintra (Intel), Stratis Viglas (University of Edinburgh)
- **ThyNVM: Enabling Software-Transparent Crash Consistency in Persistent Memory Systems** 672
Jinglei Ren (Tsinghua University), Jishen Zhao (University of California, Santa Cruz), Samira Khan (University of Virginia), Jongmoo Choi (Dankook University), Yongwei Wu (Tsinghua University), Onur Mutlu (Carnegie Mellon University)
- **Coherence Domain Restriction on Large Scale Systems** 686
Yaosheng Fu (Princeton University), Tri M. Nguyen (Princeton University), David Wentzlaff (Princeton University)
- **Efficiently Enforcing Strong Memory Ordering in GPUs** 699
Abhayendra Singh (University of Michigan), Shaizeen Aga (University of Michigan), Satish Narayanasamy (University of Michigan)

Session G2: Modeling & Characterization

- **Characterizing, Modeling, and Improving the QoE of Mobile Devices with Low Battery Level**713
Kaige Yan (University of Houston), Xingyao Zhang (University of Houston), Xin Fu (University of Houston)
- **Cross-Architecture Performance Prediction (XAPP): Using CPU Code to Predict GPU Performance**..... 725
Newsha Ardalani (University of Wisconsin, Madison), Clint Lestourgeon (University of Wisconsin, Madison), Karthikeyan Sankaralingam (University of Wisconsin, Madison), Xiaojin Zhu (University of Wisconsin, Madison)
- **A Fast and Accurate Analytical Technique to Compute the AVF of Sequential Bits in a Processor**..... 738
Steven Raasch (Intel), Arijit Biswas (Intel), Jon Stephan (Intel), Paul Racunas (Nvidia), Joel Emer (Nvidia)
- **Enabling Portable Energy Efficiency with Memory Accelerated Library** 750
Qi Guo (Carnegie Mellon University), Tze-Meng Low (Carnegie Mellon University), Nikolaos Alachiotis (Carnegie Mellon University), Berkin Akin (Carnegie Mellon University), Larry Pileggi (Carnegie Mellon University), James C. Hoe (Carnegie Mellon University), Franz Franchetti (Carnegie Mellon University)
- **Microarchitectural Implications of Event-driven Server-side Web Applications** 762
Yuhao Zhu (University of Texas at Austin), Daniel Richins (University of Texas at Austin), Matthew Halpern (University of Texas at Austin), Vijay Janapa Reddi (University of Texas at Austin)