

2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA 2015)

**Portland, Oregon, USA
13-17 June 2015**



**IEEE Catalog Number: CFP15030-POD
ISBN: 978-1-5090-0257-3**

Table of Contents

ISCA 2015 Organization Committee.....	xiii
ISCA 2015 Program Committee.....	xiv
ISCA 2015 External Review Committee.....	xvi
ISCA 2015 External Reviewers.....	xviii
ISCA 2015 Sponsors and Supporters.....	xix

Session 1: Datacenter Architectures I

- **BlueDBM: An Appliance for Big Data Analytics.....1**
Sang-Woo Jun (MIT), Ming Liu (MIT), Sungjin Lee (MIT), Jamey Hicks (Quanta Research Cambridge), John Ankcorn (Quanta Research Cambridge), Myron King (Quanta Research Cambridge), Shuotao Xu (MIT), and Arvind (MIT)
- **Towards Sustainable In-Situ Server Systems in the Big Data Era.....14**
Chao Li (Shanghai Jiao Tong University/University of Florida), Yang Hu (University of Florida), Longjun Liu (Xi'an Jiaotong University/University of Florida), Juncheng Gu (University of Florida), Mingcong Song (University of Florida), Xiaoyao Liang (Shanghai Jiao Tong University), Jingling Yuan (Wuhan University of Technology), and Tao Li (University of Florida)
- **DjiNN and Tonic: DNN as a Service and Its Implications for Future Warehouse Scale Computers.....27**
Johann Hauswald, Yiping Kang, Michael A. Laurenzano, Quan Chen, Cheng Li, Trevor Mudge, Ronald G. Dreslinski, Jason Mars, and Lingjia Tang (University of Michigan)

Session 2A: GPUs I

- **A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Flexible Data Compression with Assist Warps.....41**
Nandita Vijaykumar (CMU), Gennady Pekhimenko (CMU), Adwait Jog (Pennsylvania State University), Abhishek Bhowmick (CMU), Rachata Ausavarungnirun (CMU), Chita Das (Pennsylvania State University), Mahmut Kandemir (Pennsylvania State University), Todd C. Mowry (CMU), and Onur Mutlu (CMU)
- **Harmonia: Balancing Compute and Memory Power in High-Performance GPUs.....54**
Indrani Paul (AMD Research/Georgia), Wei Huang (AMD), Manish Arora (AMD/UCSD), and Sudhakar Yalamanchili (Georgia Tech)

Session 2B: Virtual Memory Management

- **Redundant Memory Mappings for Fast Access to Large Memories.....66**
Vasileios Karakostas (Barcelona Supercomputing Center/UPC), Jayneel Gandhi (University of Wisconsin-Madison), Furkan Ayar (Dumlupinar University), Adrián Cristal (Barcelona Supercomputing Center/UPC), Mark D. Hill (University of Wisconsin-Madison), Kathryn S. McKinley (Microsoft Research), Mario Nemirovsky (Barcelona Supercomputing Center), Michael M. Swift (University of Wisconsin-Madison), and Osman Ünsal (Barcelona Supercomputing Center)

- **Page Overlays: An Enhanced Virtual Memory Framework to Enable Fine-grained Memory Management.....79**
Vivek Seshadri (CMU), Gennady Pekhimenko (CMU), Olatunji Ruwase (Microsoft), Onur Mutlu (CMU), Phillip B. Gibbons (Intel), Michael A. Kozuch (Intel), Todd C. Mowry (CMU), and Trishul Chilimbi (Microsoft)

Session 3A: Accelerators I

- **ShiDianNao: Shifting Vision Processing Closer to the Sensor.....92**
Zidong Du (Chinese Academy of Sciences), Robert Fasthuber (EPFL), Tianshi Chen (Chinese Academy of Sciences), Paolo Ienne (EPFL), Ling Li (Chinese Academy of Sciences), Tao Luo (Chinese Academy of Sciences), Xiaobing Feng (Chinese Academy of Sciences), Yunji Chen (Chinese Academy of Sciences), and Olivier Temam (INRIA)
- **A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing.....105**
Junwhan Ahn (Seoul National University), Sungpack Hong (Oracle), Sungjoo Yoo (Seoul National University), Onur Mutlu (CMU), and Kiyoungh Choi (Seoul National University)
- **Efficient Execution of Memory Access Phases Using Dataflow Specialization.....118**
Chen-Han Ho, Sung Jin Kim, and Karthikeyan Sankaralingam (University of Wisconsin-Madison)
- **Data Reorganization in Memory Using 3D-stacked DRAM.....131**
Berkin Akin, Franz Franchetti, and James C. Hoe (CMU)

Session 3B: Performance Analysis and Tools

- **Quantitative Comparison of Hardware Transactional Memory for Blue Gene/Q, zEnterprise EC12, Intel Core, and POWER8.....144**
Takuya Nakaike (IBM), Rei Odaira (IBM), Matthew Gaudet (IBM), Maged M. Michael (IBM), and Hisanobu Tomari (University of Tokyo)
- **Profiling a warehouse-scale computer.....158**
Svilen Kanev (Harvard), Juan Pablo Darago (Universidad de Buenos Aires), Kim Hazelwood (Yahoo), Parthasarathy Ranganathan (Google), Tipp Moseley (Google), Gu-Yeon Wei (Harvard), and David Brooks (Harvard)
- **Computer Performance Microscopy with SHIM.....170**
Xi Yang (Australian National University), Stephen M. Blackburn (Australian National University), and Kathryn S. McKinley (Microsoft)
- **Flexible Software Profiling of GPU Architectures.....185**
Mark Stephenson (NVIDIA), Siva Hari (NVIDIA), Yunsup Lee (NVIDIA/UC Berkeley), Eiman Ebrahimi (NVIDIA), Daniel Johnson (NVIDIA), David Nellans (NVIDIA), Mike O'Connor (NVIDIA/UT-Austin), and Stephen W. Keckler (NVIDIA/UT-Austin)

Session 4A: DRAM Caches and Architectures

- **BEAR: Techniques for Mitigating Bandwidth Bloat in Gigascale DRAM Caches.....198**
Chiachen Chou (Georgia Tech), Aamer Jaleel (NVIDIA), and Moinuddin K. Qureshi (Georgia Tech)

- **A Fully Associative, Tagless DRAM Cache.....211**
Yongjun Lee (Sungkyunkwan University), Jongwon Kim (Sungkyunkwan University), Hakbeom Jang (Sungkyunkwan University), Hyunggyun Yang (POSTECH), Jangwoo Kim (POSTECH), Jinkyu Jeong (Sungkyunkwan University), and Jae W. Lee (Sungkyunkwan University)
- **Multiple Clone Row DRAM: A Low Latency and Area Optimized DRAM.....223**
Jungwhan Choi (KAIST), Wongyu Shin (KAIST), Jaemin Jang (KAIST), Jinwoong Suh (KAIST), Yongkee Kwon (SK Hynix), Youngsuk Moon (SK Hynix), and Lee-Sup Kim (KAIST)
- **Flexible Auto-Refresh: Enabling Scalable and Energy-Efficient DRAM Refresh Reductions.....235**
Ishwar Bhati (Oracle), Zeshan Chishti (Intel), Shih-Lien Lu (Intel), and Bruce Jacob (University of Maryland)

Session 4B: Processor Architecture I

- **Cost-Effective Speculative Scheduling in High Performance Processors.....247**
Arthur Perais (INRIA), André Seznec (INRIA), Pierre Michaud (INRIA), Andreas Sembrant (Uppsala University), and Erik Hagersten (Uppsala University)
- **LaZy Superscalar.....260**
Gorkem Asilioglu, Zhaoxiang Jin, Murat Koksal, Omkar Javeri, and Soner Onder (Michigan Tech)
- **The Load Slice Core Microarchitecture.....272**
Trevor E. Carlson (Uppsala University), Wim Heirman (Intel), Osman Allam (Ghent University), Stefanos Kaxiras (Uppsala University), and Lieven Eeckhout (Ghent University)
- **Semantic Locality and Context-based Prefetching using Reinforcement Learning.....285**
Leeor Peled (Technion/Intel), Shie Mannor (Technion), Uri Weiser (Technion), and Yoav Etsion (Technion)

Session 5: Processor Architecture II

- **Exploring the Potential of Heterogeneous Von Neumann/Dataflow Execution Models.....298**
Tony Nowatzki, Vinay Gangadhar, and Karthikeyan Sankaralingam (University of Wisconsin-Madison)
- **SHRINK: Reducing the ISA Complexity Via Instruction Recycling.....311**
Bruno Cardoso Lopes, Rafael Auler, Luiz Ramos, Edson Borin, and Rodolfo Azevedo (UNICAMP)
- **Branch Vanguard: Decomposing Branch Functionality into Prediction and Resolution Instructions...323**
Daniel McFarlin (CMU) and Craig Zilles (UIUC)

Session 6A: Memory Systems I

- **PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture....336**
Junwhan Ahn (Seoul National University), Sungjoo Yoo (Seoul National University), Onur Mutlu (CMU), Kiyoun Choi (Seoul National University)
- **SLIP: Reducing Wire Energy in the Memory Hierarchy.....349**
Subhasis Das (Stanford), Tor M. Aamodt (University of British Columbia), William J. Dally (Stanford/NVIDIA)

Session 6B: Security and Virtualization

- **CloudMonatt: an Architecture for Security Health Monitoring and Attestation of Virtual Machines in Cloud Computing.....362**
Tianwei Zhang and Ruby B. Lee (Princeton)
- **Reducing World Switches in Virtualized Environment with Flexible Cross-world Calls.....375**
Wenhao Li, Yubin Xia, Haibo Chen, Binyu Zang, and Haibin Guan (Shanghai Jiao Tong University)

Session 7A: Parallel Architectures

- **ArMOR: Defending Against Memory Consistency Model Mismatches in Heterogeneous Architectures..388**
Daniel Lustig (Princeton), Caroline Trippel (Princeton), Michael Pellauer (NVIDIA), and Margaret Martonosi (Princeton)
- **CLEAN: A Race Detector with Cleaner Semantics.....401**
Cedomir Segulja and Tarek Abdelrahman (University of Toronto)
- **MiSAR: Minimalistic Synchronization Accelerator with Resource Overflow Management.....414**
Ching-Kai Liang and Milos Prvulovic (Georgia Tech)
- **Callback: Efficient Synchronization without Invalidation with a Directory Just for Spin-Waiting.....427**
Alberto Ros (Universidad de Murcia) and Stefanos Kaxiras (Uppsala University)

Session 7B: Datacenter Architectures II

- **Thermal Time Shifting: Leveraging Phase Change Materials to Reduce Cooling Costs in Warehouse-Scale Computers.....439**
Matt Skach (University of Michigan), Manish Arora (UCSD/AMD), Chang-Hong Hsu (University of Michigan), Qi Li (UCSD), Dean Tullsen (UCSD), Lingjia Tang (University of Michigan), and Jason Mars (University of Michigan)
- **Heracles: Improving Resource Efficiency at Scale.....450**
David Lo (Stanford/Google), Liqun Cheng (Google), Rama Govindaraju (Google), Parthasarathy Ranganathan (Google), and Christos Kozyrakis (Stanford)
- **HEB: Deploying and Managing Hybrid Energy Buffers for Improving Datacenter Efficiency and Economy.....463**
Longjun Liu (Xi'an Jiaotong University), Chao Li (Shanghai Jiao Tong University), Hongbin Sun (Xi'an Jiaotong University), Yang Hu (University of Florida), Juncheng Gu (University of Florida), Tao Li (University of Florida), Jingmin Xin (Xi'an Jiaotong University), and Nanning Zheng (Xi'an Jiaotong University)
- **Architecting to Achieve a Billion Requests Per Second Throughput on a Single Key-Value Store Server Platform.....476**
Sheng Li (Intel), Hyeontaek Lim (CMU), Victor Lee (Intel), Jung Ho Ahn (Seoul National University), Anuj Kalia (CMU), Michael Kaminsky (Intel), David Andersen (CMU), Seongil O (Seoul National University), Sukhan Lee (Seoul National University), and Pradeep Dubey (Intel)

Session 8: GPUs II

- **A Variable Warp Size Architecture.....489**
Timothy G. Rogers (University of British Columbia), Daniel R. Johnson (NVIDIA), Mike O'Connor (NVIDIA/UT-Austin), and Stephen W. Keckler (NVIDIA/UT-Austin)
- **Warped-Compression: Enabling Power Efficient GPUs through Register Compression.....502**
Sangpil Lee (Yonsei University), Keunsoo Kim (Yonsei University), Gunjae Koo (USC), Hyeran Jeon (USC), Won Woo Ro (Yonsei University), and Murali Annavaram (USC)
- **CAWA: Coordinated Warp Scheduling and Cache Prioritization for Critical Warp Acceleration of GPGPU Workloads.....515**
Shin-Ying Lee, Akhil Arunkumar, and Carole-Jean Wu (Arizona State)
- **Dynamic Thread Block Launch: A Lightweight Execution Mechanism to Support Irregular Applications on GPUs.....528**
Jin Wang (Georgia Tech), Norm Rubin (NVIDIA), Albert Sidelnik (NVIDIA), and Sudhakar Yalamanchili (Georgia Tech)

Session 9A: Accelerators II

- **DynaSpAM: Dynamic Spatial Architecture Mapping using Out of Order Instruction Schedules.....541**
Feng Liu, Heejin Ahn, Stephen R. Beard, Taewook Oh, and David August (Princeton)
- **Rumba: An Online Quality Management System for Approximate Computing.....554**
Daya S Khudia, Babak Zamirai, Mehrzad Samadi, and Scott Mahlke (University of Michigan)

Session 9B: Networks and Storage

- **Manycore Network Interfaces for In-Memory Rack-Scale Computing.....567**
Alexandros Daglis (EPFL), Stanko Novakovic (EPFL), Edouard Bugnion (EPFL), Babak Falsafi (EPFL), and Boris Grot (University of Edinburgh)
- **Unified Address Translation for Memory-Mapped SSDs with FlashMap.....580**
Jian Huang (Georgia Tech), Anirudh Badam (Microsoft), Moinuddin K. Qureshi (Georgia Tech), and Karsten Schwan (Georgia Tech)

Session 10A: Security

- **FASE: Finding Amplitude-modulated Side-channel Emanations.....592**
Robert Callan, Alenka Zajic, and Milos Prvulovic (Georgia Tech)
- **Probable Cause: The Deanonymizing Effects of Approximate DRAM.....604**
Amir Rahmati, Matthew Hicks, Daniel E. Holcomb, and Kevin Fu (University of Michigan)
- **PrORAM: Dynamic Prefetcher for Oblivious RAM.....616**
Xiangyao Yu (MIT), Syed Kamran Haider (University of Connecticut), Ling Ren (MIT), Christopher Fletcher (MIT), Albert Kwon (MIT), Marten van Dijk (University of Connecticut), and Srinivas Devadas (MIT)

Session 10B: Mobile and Embedded Systems

- **MBus: An Ultra-Low Power Interconnect Bus for Next Generation Nanopower Systems.....629**
Pat Pannuto, Yoonmyung Lee, Ye-Sheng Kuo, Zhi Yoong Foo, Benjamin Kempke, Gyouho Kim, Ronald Dreslinski Jr., David Blaauw, and Prabal Dutta (University of Michigan)
- **Accelerating Asynchronous Programs through Event Sneak Peak.....642**
Gaurav Chadha, Scott Mahlke, and Satish Narayanasamy (University of Michigan)
- **VIP: Virtualizing IP Chains on Handheld Platforms.....655**
Nachiappan Chidambaram Nachiappan (Penn State), Haibo Zhang (Penn State), Jihyun Ryoo (Penn State), Niranjana Soundararajan (Intel), Anand Sivasubramaniam (Penn State), Mahmut Kandemir (Penn State), Ravishankar Iyer (Intel), and Chita R. Das (Penn State)

Session 11A: Dependable Architectures

- **FaultHound: Value-Locality-Based Soft-Fault Tolerance.....668**
Nitin, Irith Pomeranz, and T. N. Vijaykumar (Purdue)
- **COP: To Compress and Protect Main Memory.....682**
David J. Palframan, Nam Sung Kim, and Mikko H. Lipasti (University of Wisconsin-Madison)
- **Hi-fi Playback: Tolerating Position Errors in Shift Operations of Racetrack Memory.....694**
Chao Zhang (Peking University), Guangyu Sun (Peking University), Xian Zhang (Peking University), Weiqi Zhang (Peking University), Weisheng Zhao (Beihang University), Tao Wang (Peking University), Yun Liang (Peking University), Yongpan Liu (Tsinghua University), Yu Wang (Tsinghua University), and Jiwu Shu (Tsinghua University)

Session 11B: Memory Systems II

- **Stash: Have Your Scratchpad and Cache it Too.....707**
Rakesh Komuravelli, Matthew D. Sinclair, Johnathan Alsop, Muhammad Huzaifa, Maria Kotsifakou, Prakash Srivastava, Sarita V. Adve, and Vikram Adve (UIUC)
- **Coherence Protocol for Transparent Management of Scratchpad Memories in Shared Memory Manycore Architectures.....720**
Lluç Alvarez (Barcelona Supercomputing Center/UPC), Lluç Vilanova (Barcelona Supercomputing Center/UPC), Miquel Moreto (Barcelona Supercomputing Center), Marc Casas (Barcelona Supercomputing Center), Marc González (UPC), Xavier Martorell (Barcelona Supercomputing Center/UPC), Nacho Navarro (Barcelona Supercomputing Center/UPC), Eduard Ayguadé (Barcelona Supercomputing Center/UPC), and Mateo Valero (Barcelona Supercomputing Center/UPC)
- **Fusion: Design Tradeoffs in Coherent Cache Hierarchies for Accelerators.....733**
Snehasish Kumar, Arrvindh Shriraman, and Naveen Vedula (Simon Fraser University)