ISCA 2014

1. **ArchRanker: A Ranking Approach to Design Space Exploration**.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)
2. **Aladdin: A Pre-RTL, Power-Performance Accelerator Simulator Enabling Large Design Space Exploration of Customized Architectures**. Yakun Sophia Shao, Brandon Reagen, Gu-Yeon Wei, David Brooks (Harvard University)
3. **Harnessing ISA Diversity: Design of a Heterogeneous-ISA Chip Multiprocessor**. Ashish Venkat, Dean M. Tullsen (University of California, San Diego)
4. **The Direct-to-Data (D2D) Cache: Navigating the Cache Hierarchy with a Single Lookup**. Andreas Sembrant, Erik Hagersten, David Black-Schaffer (Uppsala University)
5. **The Dirty-Block Index**. 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)
6. **SC2: A Statistical Compression Cache Scheme.** Angelos Arelakis, Per Stenstrom (Chalmers University of Technology)
7. **Towards Energy Proportionality for Large-Scale Latency-Critical Workloads**. David Lo (Stanford University), Liqun Cheng (Google), Rama Govindaraju (Google), Luiz André Barroso (Google), Christos Kozyrakis (Stanford University)
8. **The CHERI capability model: Revisiting RISC in an age of risk**. 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)
9. **EOLE: Paving the Way for an Effective Implementation of Value Prediction.** Arthur Perais, André Seznec (IRISA/INRIA)
10. **Improving the Energy Efficiency of Big Cores**. 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)
11. **General-Purpose Code Acceleration with Limited-Precision Analog Computation**. 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)
12. **WebCore: Architectural Support for Mobile Web Browsing**. Yuhao Zhu, Vijay Janapa Reddi (The University of Texas at Austin)

ASPLOS 2014

13. **Uncertain< T >: A First-Order Type for Uncertain Data**. James Bornholt (Australian National Universit); Todd Mytkowicz (Microsoft Research); Kathryn S. McKinley (University of Texas at Austin/Microsoft Research)
14. **Q100: The Architecture and Design of a Database Processing Unit**. Lisa Wu (Columbia University); Andrea Lottarini (Columbia University); Timothy K. Paine (Columbia University); Martha A. Kim (Columbia University); Kenneth A. Ross (Columbia University)
15. **SI-TM: Reducing Transactional Memory Abort Rates through Snapshot Isolation**. Heiner Litz (Stanford University); David Cheriton (Stanford University); Omid Azizi (Hicamp Systems); Amin Firoozshahian (Hicamp Systems); John P. Stevenson (Stanford University)
16. **NVM Duet: Unified Working Memory and Persistent Store Architecture**. Ren-Shuo Liu (National Taiwan University, Taipei, Taiwan); De-Yu Shen (National Taiwan University, Taipei, Taiwan); Chia-Lin Yang (National Taiwan University, Taipei, Taiwan); Shun-Chih Yu (National Taiwan University, Taipei, Taiwan); Cheng-Yuan Michael Wang (Macronix International Co., Ltd., Hsinchu, Taiwan)
17. **Integrated 3D-Stacked Server Designs for Increasing Physical Density of Key-Value Stores**. Anthony Gutierrez (University of Michigan); Michael Cieslak (University of Michigan); Bharan Giridhar (University of Michigan); Ronald G. Dreslinski (University of Michigan); Luis Ceze (University of Washington); Trevor Mudge (University of Michigan)
18. **The Benefit of SMT in the Multi-Core Era: Flexibility towards Degrees of Thread-Level Parallelism**. Stijn Eyerman (Ghent University); Lieven Eeckhout (Ghent University)
19. **Locality-Oblivious Cache Organization leveraging Single-Cycle Multi-Hop NoCs**. Woo Cheol Kwon (MIT); Tushar Krishna (MIT); Li-Shiuan Peh (MIT)

20. **Ubik: Efficient Cache Sharing with Strict QoS for Latency-Critical Workloads**. Harshad Kasture (MIT); Daniel Sanchez (MIT)

HPCA 2014

21. **DASCA: Dead Write Prediction Assisted STT-RAM Cache Architecture**. Junwhan Ahn (Seoul National University), Sungjoo Yoo (Pohang University of Science and Technology), Kiyoung Choi (Seoul National University)
22. **Exploiting Thermal Energy Storage to Reduce Data Center Capital and Operating Expenses**. Wenli Zheng (The Ohio State University), Kai Ma (The Ohio State University), Xiaorui Wang (The Ohio State University)
23. **Improving Cache Performance by Exploiting Read-Write Disparity**. Samira Khan (Intel/Carnegie Mellon University), Alaa R. Alameldeen (Intel), Chris Wilkerson (Intel), Onur Mutlu (Carnegie Mellon University), Daniel A. Jimenez (Texas A&M University)
24. **Reducing the Cost of Persistence for Nonvolatile Heaps in End User Devices**. Sudarsun Kannan (Georgia Institute of Technology), Ada Gavrilovska (Georgia Institute of Technology), Karsten Schwan (Georgia Institute of Technology)
25. **Dynamic Management of TurboMode in Modern Multi-core Chips**. David Lo (Stanford), Christos Kozyrakis (Stanford)

MICRO 2013

26. **Approximate Storage in Solid-State Memories**. Adrian Sampson, Jacob Nelson (University of Washington), Karin Strauss (Microsoft Research and University of Washington), and Luis Ceze (University of Washington), http://dl.acm.org/citation.cfm?doid=2540708.2540712
27. **Decoupled Compressed Cache: Exploiting Spatial Locality for Energy-Optimized Compressed Caching**. Somayeh Sardashti and David A. Wood (University of Wisconsin-Madison), http://dl.acm.org/citation.cfm?doid=2540708.2540715
28. **Linearly Compressed Pages: A Low-Complexity, Low-Latency Main Memory Compression Framework**. Gennady Pekhimenko, Vivek Seshadri, Yoongu Kim, Hongyi Xin, Onur Mutlu (CMU), Phillip B. Gibbons, Michael A. Kozuch (Intel), and Todd C. Mowry (CMU), http://dl.acm.org/citation.cfm?doid=2540708.2540724
29. **RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization**. Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu (CMU), Phillip B. Gibbons, Michael A. Kozuch (Intel), and Todd C. Mowry (CMU), http://dl.acm.org/citation.cfm?doid=2540708.2540725
30. **Crank It Up or Dial It Down: Coordinated Multiprocessor Frequency and Folding Control**. Augusto Vega, Alper Buyuktosunoglu, Heather Hanson, Pradip Bose, and Srinivasan Ramani (IBM), http://dl.acm.org/citation.cfm?doid=2540708.2540727
31. **Wavelength Stealing: An Opportunistic Approach to Channel Sharing in Multi-chip Photonic Interconnects**. Arslan Zulfiqar (University of Wisconsin - Madison), Pranay Koka, Herb Schwetman (Oracle Labs), Mikko Lipasti (University of Wisconsin - Madison), Xuezhe Zheng, and Ashok V. Krishnamoorthy (Oracle Labs), http://dl.acm.org/citation.cfm?doid=2540708.2540728
32. **Linearizing Irregular Memory Accesses for Improved Correlated Prefetching**. Akanksha Jain and Calvin Lin (University of Texas at Austin), http://dl.acm.org/citation.cfm?doid=2540708.2540730
33. **Imbalanced Cache Partitioning for Balanced Data-Parallel Programs**. Abhisek Pan and Vijay S. Pai (Purdue University), http://dl.acm.org/citation.cfm?doid=2540708.2540734
34. **Kiln: Closing the Performance Gap Between Systems With and Without Persistence Support**. Jishen Zhao (Pennsylvania State University), Sheng Li (Hewlett-Packard Labs), Doe Hyun Yoon (IBM Thomas J. Watson Research Center), Yuan Xie (Pennsylvania State University/AMD Research), and Norman P. Jouppi (Google), http://dl.acm.org/citation.cfm?doid=2540708.2540744
35. **Meet the Walkers: Accelerating Index Traversals for In-Memory Databases**. Onur Kocberber (EPFL), Boris Grot (University of Edinburgh), Javier Picorel, Babak Falsafi (EPFL), Kevin Lim (HP Labs), and Parthasarathy Ranganathan (Google), http://dl.acm.org/citation.cfm?doid=2540708.2540748