

2014 47th Annual IEEE/ACM International Symposium on Microarchitecture

(MICRO 2014)

**Cambridge, United Kingdom
13-17 December 2014**



**IEEE Catalog Number: CFP14071-POD
ISBN: 978-1-4799-6999-9**

2014 47th Annual IEEE/ACM International Symposium on Microarchitecture

MICRO 2014

Table of Contents

Message from the General Chair	x
Message from the Program Co-Chairs	xi
Organizing Committee	xiii
Program Committee	xiv
External Reviewer Committee	xvi
External Reviewers	xix
Keynote Abstracts	xx

Session 1A: Stacked DRAM

CAMEO: A Two-Level Memory Organization with Capacity of Main Memory and Flexibility of Hardware-Managed Cache	1
<i>Chia Chen Chou, Aamer Jaleel, and Moinuddin K. Qureshi</i>	
Transparent Hardware Management of Stacked DRAM as Part of Memory	13
<i>Jaewoong Sim, Alaa R. Alameldeen, Zeshan Chishti, Chris Wilkerson, and Hyesoon Kim</i>	
Unison Cache: A Scalable and Effective Die-Stacked DRAM Cache	25
<i>Djordje Jevdjic, Gabriel H. Loh, Cansu Kaynak, and Babak Falsafi</i>	
Bi-Modal DRAM Cache: Improving Hit Rate, Hit Latency and Bandwidth	38
<i>Nagendra Gulur, Mahesh Mehendale, R. Manikantan, and R. Govindarajan</i>	
Citadel: Efficiently Protecting Stacked Memory from Large Granularity Failures	51
<i>Prashant J. Nair, David A. Roberts, and Moinuddin K. Qureshi</i>	

Session 1B: GPGPU and Data Parallel Architectures

Locality-Aware Mapping of Nested Parallel Patterns on GPUs	63
<i>Hyoukjoong Lee, Kevin J. Brown, Arvind K. Sujeeth, Tiark Rompf, and Kunle Olukotun</i>	
Accelerating Irregular Algorithms on GPGPUs Using Fine-Grain Hardware Worklists	75
<i>Ji Yun Kim and Christopher Batten</i>	

PORPLE: An Extensible Optimizer for Portable Data Placement on GPU	88
<i>Guoyang Chen, Bo Wu, Dong Li, and Xipeng Shen</i>	
Exploring the Design Space of SPMD Divergence Management on Data-Parallel Architectures	101
<i>Yunsup Lee, Vinod Grover, Ronny Krashinsky, Mark Stephenson, Stephen W. Keckler, and Krste Asanović</i>	
Managing GPU Concurrency in Heterogeneous Architectures	114
<i>Onur Kayiran, Nachiappan Chidambaram Nachiappan, Adwait Jog, Rachata Ausavarungnirun, Mahmut T. Kandemir, Gabriel H. Loh, Onur Mutlu, and Chita R. Das</i>	
Session 2A: Memory Systems, Scheduling, and Optimization	
Load Value Approximation	127
<i>Joshua San Miguel, Mario Badr, and Natalie Enright Jerger</i>	
Arbitrary Modulus Indexing	140
<i>Jeffrey R. Diamond, Donald S. Fussell, and Stephen W. Keckler</i>	
FIRM: Fair and High-Performance Memory Control for Persistent Memory Systems	153
<i>Jishen Zhao, Onur Mutlu, and Yuan Xie</i>	
Short-Circuiting Memory Traffic in Handheld Platforms	166
<i>Praveen Yedlapalli, Nachiappan Chidambaram Nachiappan, Niranjan Soundararajan, Anand Sivasubramaniam, Mahmut T. Kandemir, and Chita R. Das</i>	
Efficient Memory Virtualization: Reducing Dimensionality of Nested Page Walks	178
<i>Jayneel Gandhi, Arkaprava Basu, Mark D. Hill, and Michael M. Swift</i>	
Session 2B: Security	
Iso-X: A Flexible Architecture for Hardware-Managed Isolated Execution	190
<i>Dmitry Evtuyshkin, Jesse Elwell, Meltem Ozsoy, Dmitry Ponomarev, Nael Abu Ghazaleh, and Ryan Riley</i>	
Random Fill Cache Architecture	203
<i>Fangfei Liu and Ruby B. Lee</i>	
CC-Hunter: Uncovering Covert Timing Channels on Shared Processor Hardware	216
<i>Jie Chen and Guru Venkataramani</i>	
Continuous, Low Overhead, Run-Time Validation of Program Executions	229
<i>Erdem Aktas, Furat Afram, and Kanad Ghose</i>	
A Practical Methodology for Measuring the Side-Channel Signal Available to the Attacker for Instruction-Level Events	242
<i>Robert Callan, Alenka Zajić, and Milos Prvulovic</i>	

Session 3A: Methodology, Modeling, and Tools

RpStacks: Fast and Accurate Processor Design Space Exploration Using Representative Stall-Event Stacks	255
<i>Jaewon Lee, Hanhwi Jang, and Jangwoo Kim</i>	
GPUMech: GPU Performance Modeling Technique Based on Interval Analysis	268
<i>Jen-Cheng Huang, Joo Hwan Lee, Hyesoon Kim, and Hsien-Hsin S. Lee</i>	
PyMTL: A Unified Framework for Vertically Integrated Computer Architecture Research	280
<i>Derek Lockhart, Gary Zibrat, and Christopher Batten</i>	

Session 3B: Reliability and Fault Tolerance

Calculating Architectural Vulnerability Factors for Spatial Multi-Bit Transient Faults	293
<i>Mark Wilkening, Vilas Sridharan, Si Li, Fritz Previlon, Sudhanva Gurumurthi, and David R. Kaeli</i>	
Using ECC Feedback to Guide Voltage Speculation in Low-Voltage Processors	306
<i>Anys Bacha and Radu Teodorescu</i>	
Harnessing Soft Computations for Low-Budget Fault Tolerance	319
<i>Daya Shanker Khudia and Scott Mahlke</i>	

Session 4A: TLB and Cache Optimization

Skewed Compressed Caches	331
<i>Somayeh Sardashti, André Seznec, and David A. Wood</i>	
Adaptive Cache Management for Energy-Efficient GPU Computing	343
<i>Xuhao Chen, Li-Wen Chang, Christopher I. Rodrigues, Jie Lv, Zhiying Wang, and Wen-Mei Hwu</i>	
Futility Scaling: High-Associativity Cache Partitioning	356
<i>Ruisheng Wang and Lizhong Chen</i>	

Session 4B: Managing Voltage and Time

Voltage Noise in Multi-Core Processors: Empirical Characterization and Optimization Opportunities	368
<i>Ramon Bertran, Alper Buyuktosunoglu, Pradip Bose, Timothy J. Slegel, Gerard Salem, Sean Carey, Richard F. Rizzolo, and Thomas Strach</i>	
Enabling Realistic Fine-Grain Voltage Scaling with Reconfigurable Power Distribution Networks	381
<i>Waclaw Godycki, Christopher Torng, Ivan Bukreyev, Alyssa Apsel, and Christopher Batten</i>	
Micro-Sliced Virtual Processors to Hide the Effect of Discontinuous CPU Availability for Consolidated Systems	394
<i>Jeongseob Ahn, Chang Hyun Park, and Jaehyuk Huh</i>	

Session 5A: Energy-Efficient Computation

SMiTe: Precise QoS Prediction on Real-System SMT Processors to Improve Utilization in Warehouse Scale Computers	406
<i>Yunqi Zhang, Michael A. Laurenzano, Jason Mars, and Lingjia Tang</i>	
A Front-End Execution Architecture for High Energy Efficiency	419
<i>Ryota Shioya, Masahiro Goshima, and Hideki Ando</i>	
Execution Drafting: Energy Efficiency through Computation Deduplication	432
<i>Michael Mckeown, Jonathan Balkind, and David Wentzloff</i>	
PPEP: Online Performance, Power, and Energy Prediction Framework and DVFS Space Exploration	445
<i>Bo Su, Junli Gu, Li Shen, Wei Huang, Joseph L. Greathouse, and Zhiying Wang</i>	

Session 5B: Interconnects

NoC Architectures for Silicon Interposer Systems: Why Pay for more Wires when you Can Get them (from your interposer) for Free?	458
<i>Natalie Enright Jerger, Ajaykumar Kannan, Zimo Li, and Gabriel H. Loh</i>	
Hi-Rise: A High-Radix Switch for 3D Integration with Single-Cycle Arbitration	471
<i>Supreet Jeloka, Reetuparna Das, Ronald G. Dreslinski, Trevor Mudge, and David Blaauw</i>	
Multi-GPU System Design with Memory Networks	484
<i>Gwangsun Kim, Minseok Lee, Jiyun Jeong, and John Kim</i>	
Dodec: Random-Link, Low-Radix On-Chip Networks	496
<i>Haofan Yang, Jyoti Tripathi, Natalie Enright Jerger, and Dan Gibson</i>	

Session 6A: Branch Prediction and Prefetching

Wormhole: Wisely Predicting Multidimensional Branches	509
<i>Jorge Albericio, Joshua San Miguel, Natalie Enright Jerger, and Andreas Moshovos</i>	
Bias-Free Branch Predictor	521
<i>Dibakar Gope and Mikko H. Lipasti</i>	
Loop-Aware Memory Prefetching Using Code Block Working Sets	533
<i>Adi Fuchs, Shie Mannor, Uri Weiser, and Yoav Etsion</i>	
BuMP: Bulk Memory Access Prediction and Streaming	545
<i>Stavros Volos, Javier Picorel, Babak Falsafi, and Boris Grot</i>	

Session 6B: Compilation and Code Generation

Protean Code: Achieving Near-Free Online Code Transformations for Warehouse Scale Computers	558
<i>Michael A. Laurenzano, Yunqi Zhang, Lingjia Tang, and Jason Mars</i>	
Compiler Support for Optimizing Memory Bank-Level Parallelism	571
<i>Wei Ding, Diana Guttman, and Mahmut Kandemir</i>	
Architectural Specialization for Inter-Iteration Loop Dependence Patterns	583
<i>Shreesh Srinath, Berkin Ilbeyi, Mingxing Tan, Gai Liu, Zhiru Zhang, and Christopher Batten</i>	
Specializing Compiler Optimizations through Programmable Composition for Dense Matrix Computations	596
<i>Qing Yi, Qian Wang, and Huimin Cui</i>	

Session 7: Best Paper Nominees

DaDianNao: A Machine-Learning Supercomputer	609
<i>Yunji Chen, Tao Luo, Shaoli Liu, Shijin Zhang, Liqiang He, Jia Wang, Ling Li, Tianshi Chen, Zhiwei Xu, Ninghui Sun, and Olivier Temam</i>	
B-Fetch: Branch Prediction Directed Prefetching for Chip-Multiprocessors	623
<i>David Kadjo, Jinchun Kim, Prabal Sharma, Reena Panda, Paul Gratz, and Daniel Jimenez</i>	
Pipe Check: Specifying and Verifying Microarchitectural Enforcement of Memory Consistency Models	635
<i>Daniel Lustig, Michael Pellauer, and Margaret Martonosi</i>	
Equalizer: Dynamic Tuning of GPU Resources for Efficient Execution	647
<i>Ankit Sethia and Scott Mahlke</i>	
COMP: Compiler Optimizations for Manycore Processors	659
<i>Linhai Song, Min Feng, Nishkam Ravi, Yi Yang, and Srimat Chakradhar</i>	
Author Index	672