

# **2012 45th Annual IEEE/ACM International Symposium on Microarchitecture**

**(MICRO 2012)**

**Vancouver, British Columbia, Canada  
1 – 5 December 2012**



**IEEE Catalog Number: CFP12071-PRT  
ISBN: 978-1-49; ;/3926/6**

# 2012 IEEE/ACM 45th Annual International Symposium on Microarchitecture

## MICRO 2012

### Table of Contents

Message from the General Chair.....	ix
Message from the Program Chair.....	x
Organizing Committee.....	xiii
Program Committee.....	xiv
External Review Committee.....	xv
Reviewers.....	xvii

---

#### Session IA - Memory Systems I

FPB: Fine-grained Power Budgeting to Improve Write Throughput of Multi-level Cell Phase Change Memory .....	1
<i>Lei Jiang, Youtao Zhang, Bruce R. Childers, and Jun Yang</i>	
Leveraging Heterogeneity in DRAM Main Memories to Accelerate Critical Word Access .....	13
<i>Niladrish Chatterjee, Manjunath Shevgoor, Rajeev Balasubramonian, Al Davis, Zhen Fang, Ramesh Illikkal, and Ravi Iyer</i>	
Transactional Memory Architecture and Implementation for IBM System Z .....	25
<i>Christian Jacobi, Timothy Slegel, and Dan Greiner</i>	

#### Session IB - Fault Tolerance

Warped-DMR: Light-weight Error Detection for GPGPU .....	37
<i>Hyeran Jeon and Murali Annavaram</i>	
The Performance Vulnerability of Architectural and Non-architectural Arrays to Permanent Faults .....	48
<i>Damien Hardy, Isidoros Sideris, Nikolas Ladas, and Yiannakis Sazeides</i>	
NoCAAlert: An On-Line and Real-Time Fault Detection Mechanism for Network-on-Chip Architectures .....	60
<i>Andreas Prodromou, Andreas Panteli, Chrysostomos Nicopoulos, and Yiannakis Sazeides</i>	

## Session IIA - GPUs and SIMD

Cache-Conscious Wavefront Scheduling .....	72
<i>Timothy G. Rogers, Mike O'Connor, and Tor M. Aamodt</i>	
Libra: Tailoring SIMD Execution Using Heterogeneous Hardware and Dynamic Configurability .....	84
<i>Yongjun Park, Jason Jong Kyu Park, Hyunchul Park, and Scott Mahlke</i>	
Unifying Primary Cache, Scratch, and Register File Memories in a Throughput Processor .....	96
<i>Mark Gebhart, Stephen W. Keckler, Brucek Khailany, Ronny Krashinsky, and William J. Dally</i>	
Kernel Weaver: Automatically Fusing Database Primitives for Efficient GPU Computation .....	107
<i>Haicheng Wu, Gregory Diamos, Srihari Cadambi, and Sudhakar Yalamanchili</i>	

## Session IIB - Energy I

KnightShift: Scaling the Energy Proportionality Wall through Server-Level Heterogeneity .....	119
<i>Daniel Wong and Murali Annavaram</i>	
Rethinking DRAM Power Modes for Energy Proportionality .....	131
<i>Krishna T. Malladi, Ian Shaeffer, Liji Gopalakrishnan, David Lo, Benjamin C. Lee, and Mark Horowitz</i>	
CoScale: Coordinating CPU and Memory System DVFS in Server Systems .....	143
<i>Qingyuan Deng, David Meisner, Abhishek Bhattacharjee, Thomas F. Wenisch, and Ricardo Bianchini</i>	
Predicting Performance Impact of DVFS for Realistic Memory Systems .....	155
<i>Rustam Miftakhutdinov, Eiman Ebrahimi, and Yale N. Patt</i>	

## Session IIIA – Big Data

Vector Extensions for Decision Support DBMS Acceleration .....	166
<i>Timothy Hayes, Oscar Palomar, Osman Unsal, Adrian Cristal, and Mateo Valero</i>	
NOC-Out: Microarchitecting a Scale-Out Processor .....	177
<i>Pejman Lotfi-Kamran, Boris Grot, and Babak Falsafi</i>	
SLICC: Self-Assembly of Instruction Cache Collectives for OLTP Workloads .....	188
<i>Islam Atta, Pınar Tözün, Anastasia Ailamaki, and Andreas Moshovos</i>	

## Session IIIB – Energy II

Systematic Energy Characterization of CMP/SMT Processor Systems via Automated Micro-Benchmarks .....	199
<i>Ramon Bertran, Alper Buyuktosunoglu, Meeta S. Gupta, Marc Gonzalez, and Pradip Bose</i>	

AUDIT: Stress Testing the Automatic Way .....	212
<i>Youngtaek Kim, Lizy Kurian John, Sanjay Pant, Srilatha Manne, Michael Schulte, W. Lloyd Bircher, and Madhu S. Sibi Govindan</i>	
Accurate Fine-Grained Processor Power Proxies .....	224
<i>Wei Huang, Charles Lefurgy, William Kuk, Alper Buyuktosunoglu, Michael Floyd, Karthick Rajamani, Malcolm Allen-Ware, and Bishop Brock</i>	
<b>Session IVA – Memory Systems II</b>	
Fundamental Latency Trade-off in Architecting DRAM Caches: Outperforming Impractical SRAM-Tags with a Simple and Practical Design .....	235
<i>Moinuddin K. Qureshi and Gabe H. Loh</i>	
A Mostly-Clean DRAM Cache for Effective Hit Speculation and Self-Balancing Dispatch .....	247
<i>Jaewoong Sim, Gabriel H. Loh, Hyesoon Kim, Mike O'Connor, and Mithuna Thottethodi</i>	
CoLT: Coalesced Large-Reach TLBs .....	258
<i>Binh Pham, Viswanathan Vaidyanathan, Amer Jaleel, and Abhishek Bhattacharjee</i>	
<b>Session IVB – Interconnects</b>	
NoRD: Node-Router Decoupling for Effective Power-gating of On-Chip Routers .....	270
<i>Lizhong Chen and Timothy M. Pinkston</i>	
Dynamic Reconfiguration of 3D Photonic Networks-on-Chip for Maximizing Performance and Improving Fault Tolerance .....	282
<i>Randy Morris, Avinash Karanth Kodi, and Ahmed Louri</i>	
Addressing End-to-End Memory Access Latency in NoC-Based Multicores .....	294
<i>Akbar Sharifi, Emre Kultursay, Mahmut Kandemir, and Chita R. Das</i>	
<b>Session VA – Core Design</b>	
MorphCore: An Energy-Efficient Microarchitecture for High Performance ILP and High Throughput TLP .....	305
<i>Khubaib, M. Aater Suleman, Milad Hashemi, Chris Wilkerson, and Yale N. Patt</i>	
Composite Cores: Pushing Heterogeneity Into a Core .....	317
<i>Andrew Lukefahr, Shruti Padmanabha, Reetuparna Das, Faissal M. Sleiman, Ronald Dreslinski, Thomas F. Wensich, and Scott Mahlke</i>	
Control-Flow Decoupling .....	329
<i>Rami Sheikh, James Tuck, and Eric Rotenberg</i>	
<b>Session VB – Coherence and Consistency</b>	
Spatiotemporal Coherence Tracking .....	341
<i>Mohammad Alisafae</i>	

Predicting Coherence Communication by Tracking Synchronization Points at Run Time .....	351
<i>Socrates Demetriades and Sangyeun Cho</i>	
Vulcan: Hardware Support for Detecting Sequential Consistency Violations Dynamically .....	363
<i>Abdullah Muzahid, Shanxiang Qi, and Josep Torrellas</i>	
<b>Session VIA – Caching</b>	
Amoeba-Cache: Adaptive Blocks for Eliminating Waste in the Memory Hierarchy .....	376
<i>Snehasish Kumar, Hongzhou Zhao, Arrvindh Shriraman, Eric Matthews, Sandhya Dwarkadas, and Lesley Shannon</i>	
Improving Cache Management Policies Using Dynamic Reuse Distances .....	389
<i>Nam Duong, Dali Zhao, Taesu Kim, Rosario Cammarota, Mateo Valero, and Alexander V. Veidenbaum</i>	
<b>Session VIB – Modeling and Partitioning</b>	
Kernel Partitioning of Streaming Applications: A Statistical Approach to an NP-complete Problem .....	401
<i>Petar Radojković, Paul M. Carpenter, Miquel Moretó, Alex Ramirez, and Francisco J. Cazorla</i>	
Inferred Models for Dynamic and Sparse Hardware-Software Spaces .....	413
<i>Weidan Wu and Benjamin C. Lee</i>	
<b>Session VIIA – Dynamic Optimization and Parallelization</b>	
SMARQ: Software-Managed Alias Register Queue for Dynamic Optimizations .....	425
<i>Cheng Wang, Youfeng Wu, Hongbo Rong, and Hyunchul Park</i>	
Profiling Data-Dependence to Assist Parallelization: Framework, Scope, and Optimization .....	437
<i>Alain Ketterlin and Philippe Claus</i>	
<b>Session VIIB – Accelerators</b>	
Neural Acceleration for General-Purpose Approximate Programs .....	449
<i>Hadi Esmaeilzadeh, Adrian Sampson, Luis Ceze, and Doug Burger</i>	
Designing a Programmable Wire-Speed Regular-Expression Matching Accelerator .....	461
<i>Jan Van Lunteren, Christoph Hagleitner, Timothy Heil, Giora Biran, Uzi Shvadron, and Kubilay Atas</i>	
<b>Author Index</b> .....	473