

2011 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2011)

**Porto Alegre, Brazil
3-7 December 2011**



**IEEE Catalog Number: CFP11071-POD
ISBN: 978-1-5090-6605-6**

**Copyright © 2011, Association for Computing Machinery (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: | CFP11071-POD |
| ISBN (Print-On-Demand): | 978-1-5090-6605-6 |
| ISBN (Online): | 978-1-4503-1053-6 |
| 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

| | |
|---|-----------|
| Keynotes | i |
| Future Architectures will Incorporate HPUs | i |
| <i>Doug Burger</i> | |
| Life After Dennard and How I Learned to Love the Picojoule | ii |
| <i>Steve Keckler</i> | |
| Race to Exascale: Challenges and Opportunities | iii |
| <i>Avinash Sodani</i> | |
| Session I: Best Paper Candidates A | 1 |
| Active Management of Timing Guardband to Save Energy in POWER7 | 1 |
| <i>Charles R. Lefurgy, Alan J. Drake, Michael S. Floyd, Malcolm S. Allen-Ware, Bishop Brock, Jose A. Tierno, and John B. Carter</i> | |
| Bundled Execution of Recurring Traces for Energy-Efficient General Purpose Processing | 12 |
| <i>Shantanu Gupta, Shuguang Feng, Amin Ansari, Scott A. Mahlke, and David I. August</i> | |
| Minimalist Open-page: A DRAM Page-mode Scheduling Policy for the Many-core Era | 24 |
| <i>Dimitris Kaseridis, Jeffrey Stuecheli, and Lizy Kurian John</i> | |
| Session II: Best Paper Candidates B | 36 |
| The NoX Router | 36 |
| <i>Mitchell Hayenga and Mikko H. Lipasti</i> | |
| A Systematic Methodology to Develop Resilient Cache Coherence Protocols | 47 |
| <i>Konstantinos Aisopos and Li-Shiuan Peh</i> | |
| Dataflow Execution of Sequential Imperative Programs on Multicore Architectures | 59 |
| <i>Gagan Gupta and Gurindar S. Sohi</i> | |
| Session III-A: NoCs | 71 |
| Towards the Ideal On-chip Fabric for 1-to-Many and Many-to-1 Communication | 71 |
| <i>Tushar Krishna, Li-Shiuan Peh, Bradford M. Beckmann, and Steven K. Reinhardt</i> | |
| Packet Chaining: Efficient Single-Cycle Allocation for On-Chip Networks | 83 |
| <i>George Michelogiannakis, Nan Jiang, Daniel Becker, and William J. Dally</i> | |

| | |
|--|------------|
| Resilient Microring Resonator Based Photonic Networks | 95 |
| <i>Christopher J. Nitta, Matthew K. Farrens, and Venkatesh Akella</i> | |
| FeatherWeight: Low-cost Optical Arbitration with QoS Support | 105 |
| <i>Yan Pan, John Kim, and Gokhan Memik</i> | |
| Session III-B: Speculation | 117 |
| <hr/> | |
| A New Case for the TAGE Branch Predictor | 117 |
| <i>André Seznec</i> | |
| Identifying and Predicting Timing-Critical Instructions to Boost Timing Speculation | 128 |
| <i>Jing Xin and Russ Joseph</i> | |
| Idempotent Processor Architecture | 140 |
| <i>Marc de Kruijf and Karthikeyan Sankaralingam</i> | |
| Proactive Instruction Fetch | 152 |
| <i>Michael Ferdman, Cansu Kaynak, and Babak Falsafi</i> | |
| Session IV-A: Energy Efficiency | 163 |
| <hr/> | |
| QSCORES: Trading Dark Silicon for Scalable Energy Efficiency with Quasi-Specific Cores | 163 |
| <i>Ganesh Venkatesh, Jack Sampson, Nathan Goulding-Hotta, Sravanthi Kota Venkata, Michael Bedford Taylor, and Steven Swanson</i> | |
| Pack & Cap: Adaptive DVFS and Thread Packing Under Power Caps | 175 |
| <i>Ryan Cochran, Can Hankendi, Ayse K. Coskun, and Sherief Reda</i> | |
| Preventing PCM Banks from Seizing Too Much Power | 186 |
| <i>Andrew Hay, Karin Strauss, Timothy Sherwood, Gabriel H. Loh, and Doug Burger</i> | |
| CRAM: Coded Registers for Amplified Multiporting | 196 |
| <i>Vignyan Reddy Kothinti Naresh, David J. Palfreman, and Mikko H. Lipasti</i> | |
| Session IV-B: Multi-Core Programmability | 206 |
| <hr/> | |
| ATDetector: Improving the Accuracy of a Commercial Data Race Detector by Identifying Address Transfer | 206 |
| <i>Jiaqi Zhang, Weiwei Xiong, Yang Liu, Soyeon Park, Yuanyuan Zhou, and Zhiqiang Ma</i> | |
| CoreRacer: A Practical Memory Race Recorder for Multicore x86 TSO Processors | 216 |
| <i>Gilles Pokam, Cristiano Pereira, Shiliang Hu, Ali-Reza Adl-Tabatabai, Justin Gottschlich, Jungwoo Ha, and Youfeng Wu</i> | |
| Manager-Client Pairing: A Framework for Implementing Coherence Hierarchies | 226 |
| <i>Jesse G. Beu, Michael C. Rosier, and Thomas M. Conte</i> | |
| TransCom: Transforming Stream Communication for Load Balance and Efficiency in Networks-on-Chip . . | 237 |
| <i>Ahmed H. Abdel-Gawad and Mithuna Thottethodi</i> | |
| Session V: Datacenters and Virtualization | 248 |
| <hr/> | |
| Bubble-Up: Increasing Utilization in Modern Warehouse Scale Computers via Sensible Co-locations | 248 |
| <i>Jason Mars, Lingjia Tang, Robert Hundt, Kevin Skadron, and Mary Lou Soffa</i> | |

| | |
|---|------------|
| System-Level Integrated Server Architectures for Scale-out Datacenters | 260 |
| <i>Sheng Li, Kevin Lim, Paolo Faraboschi, Jichuan Chang, Parthasarathy Ranganathan, and Norman P. Jouppi</i> | |
| Architectural Support for Secure Virtualization under a Vulnerable Hypervisor | 272 |
| <i>Seongwook Jin, Jeongseob Ahn, Sanghoon Cha, and Jaehyuk Huh</i> | |
| Session VI-A: Exploiting Parallelism | 284 |
| <hr/> | |
| Complementing User-Level Coarse-Grain Parallelism with Implicit Speculative Parallelism | 284 |
| <i>Nikolas Ioannou and Marcelo Cintra</i> | |
| Hardware Transactional Memory for GPU Architectures | 296 |
| <i>Wilson W. L. Fung, Inderpreet Singh, Andrew Brownsword, and Tor M. Aamodt</i> | |
| Improving GPU Performance via Large Warps and Two-Level Warp Scheduling | 308 |
| <i>Veynu Narasiman, Michael Shebanow, Chang Joo Lee, Rustam Miftakhutdinov, Onur Mutlu, and Yale N. Patt</i> | |
| Session VI-B: Memory Technologies | 318 |
| <hr/> | |
| Pay-As-You-Go: Low-Overhead Hard-Error Correction for Phase Change Memories | 318 |
| <i>Moinuddin K. Qureshi</i> | |
| Multi Retention Level STT-RAM Cache Designs with a Dynamic Refresh Scheme | 329 |
| <i>Zhenyu Sun, Xiuyuan Bi, Hai (Helen) Li, Weng-Fai Wong, Zhong-Liang Ong, Xiaochun Zhu, and Wenqing Wu</i> | |
| A Resistive TCAM Accelerator for Data-Intensive Computing | 339 |
| <i>Qing Guo, Xiaochen Guo, Yuxin Bai, and Engin Ipek</i> | |
| Session VII-A: Memory | 351 |
| <hr/> | |
| A Register-file Approach for Row Buffer Caches in Die-stacked DRAMs | 351 |
| <i>Gabriel H. Loh</i> | |
| Parallel Application Memory Scheduling | 362 |
| <i>Eiman Ebrahimi, Rustam Miftakhutdinov, Chris Fallin, Chang Joo Lee, José A. Joao, Onur Mutlu, and Yale N. Patt</i> | |
| Reducing Memory Interference in Multicore Systems via Application-Aware Memory Channel Partitioning | 374 |
| <i>Sai Prashanth Muralidhara, Lavanya Subramanian, Onur Mutlu, Mahmut Kandemir, and Thomas Moscibroda</i> | |
| Session VII-B: Validation and Reliability | 386 |
| <hr/> | |
| Accelerating Microprocessor Silicon Validation by Exposing ISA Diversity | 386 |
| <i>Nikos Foutris, Dimitris Gizopoulos, Mihalis Psarakis, Xavier Vera, and Antonio Gonzalez</i> | |
| Encore: Low-Cost, Fine-Grained Transient Fault Recovery | 398 |
| <i>Shuguang Feng, Shantanu Gupta, Amin Ansari, Scott A. Mahlke, and David I. August</i> | |
| Formally Enhanced Runtime Verification to Ensure NoC Functional Correctness | 410 |
| <i>Ritesh Parikh and Valeria Bertacco</i> | |

| | |
|---|------------|
| Session VIII: Caches | 420 |
| Residue Cache: a Low-Energy Low-Area L2 Cache Architecture via Compression and Partial Hits | 420 |
| <i>Soontae Kim, Jesung Kim, Jongmin Lee, and Seokin Hong</i> | |
| SHiP: Signature-based Hit Predictor for High Performance Caching | 430 |
| <i>Carole-Jean Wu, Amer Jaleel, Will Hasenplaugh, Margaret Martonosi, Simon C., Jr. Steely, and Joel Emer</i> | |
| PACMan: Prefetch-Aware Cache Management for High Performance Caching | 442 |
| <i>Carole-Jean Wu, Amer Jaleel, Margaret Martonosi, Simon C., Jr. Steely, and Joel Emer</i> | |
| Efficiently Enabling Conventional Block Sizes for Very Large Die-stacked DRAM Caches | 454 |
| <i>Gabriel H. Loh and Mark D. Hill</i> | |
| Session IX: Compiler Support | 465 |
| A Compile-Time Managed Multi-Level Register File Hierarchy | 465 |
| <i>Mark Gebhart, Stephen W. Keckler, and William J. Dally</i> | |
| SIMD Re-Convergence At Thread Frontiers | 477 |
| <i>Gregory Diamos, Benjamin Ashbaugh, Subramaniam Maiyuran, Andrew Kerr, Haicheng Wu, and Sudhakar Yalamanchili</i> | |
| A Data Layout Optimization Framework for NUCA-Based Multicores | 489 |
| <i>Yuanrui Zhang, Wei Ding, Mahmut Kandemir, Jun Liu, and Ohyoung Jang</i> | |
| Author Index | 501 |