

# **2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture**

**(MICRO 2010)**

**Atlanta, Georgia, USA  
4 – 8 December 2010**



**IEEE Catalog Number: CFP10071-PRT  
ISBN: 978-1-4244-9071-4**

# 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture

## *MICRO 2010*

# Table of Contents

<b>Message from the General Co-Chairs</b> .....	x
<b>Message from the Program Chair</b> .....	xi
<b>Organizing Committee</b> .....	xii
<b>Program Committee</b> .....	xiii
<b>MICRO 2010 Reviewers</b> .....	xiv

---

### **Session 1: Transactional Systems**

Scalable Speculative Parallelization on Commodity Clusters .....	3
<i>Hanjun Kim, Arun Raman, Feng Liu, Jae W. Lee, and David I. August</i>	
Hardware Support for Relaxed Concurrency Control in Transactional Memory .....	15
<i>Utku Aydonat and Tarek S. Abdelrahman</i>	
A Dynamically Adaptable Hardware Transactional Memory .....	27
<i>Marc Lupon, Grigorios Magklis, and Antonio González</i>	
ASF: AMD64 Extension for Lock-Free Data Structures and Transactional Memory .....	39
<i>Jaewoong Chung, Luke Yen, Stephan Diestelhorst, Martin Pohlack, Michael Hohmuth, David Christie, and Dan Grossman</i>	

### **Session 2A: Scheduling**

Memory Latency Reduction via Thread Throttling .....	53
<i>Hsiang-Yun Cheng, Chung-Hsiang Lin, Jian Li, and Chia-Lin Yang</i>	
Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior .....	65
<i>Yoongu Kim, Michael Papamichael, Onur Mutlu, and Mor Harchol-Balter</i>	

Voltage Smoothing: Characterizing and Mitigating Voltage Noise in Production Processors via Software-Guided Thread Scheduling .....	77
<i>Vijay Janapa Reddi, Svilen Kanev, Wonyoung Kim, Simone Campanoni, Michael D. Smith, Gu-Yeon Wei, and David Brooks</i>	
Task Superscalar: An Out-of-Order Task Pipeline .....	89
<i>Yoav Etsion, Felipe Cabarcas, Alejandro Rico, Alex Ramirez, Rosa M. Badia, Eduard Ayguade, Jesus Labarta, and Mateo Valero</i>	
<b>Session 2B: Reliability/Scheduling</b>	
Combating Aging with the Colt Duty Cycle Equalizer .....	103
<i>Erika Gunadi, Abhisek A. Sinkar, Nam Sung Kim, and Mikko H. Lipasti</i>	
SAFER: Stuck-At-Fault Error Recovery for Memories .....	115
<i>Nak Hee Seong, Dong Hyuk Woo, Vijayalakshmi Srinivasan, Jude A. Rivers, and Hsien-Hsin S. Lee</i>	
AVF Stressmark: Towards an Automated Methodology for Bounding the Worst-Case Vulnerability to Soft Errors .....	125
<i>Arun Arvind Nair, Lizy Kurian John, and Lieven Eeckhout</i>	
Flexible and Efficient Instruction-Grained Run-Time Monitoring Using On-Chip Reconfigurable Fabric .....	137
<i>Daniel Y. Deng, Daniel Lo, Greg Malysa, Skyler Schneider, and G. Edward Suh</i>	
<b>Session 3A: Caching</b>	
Achieving Non-Inclusive Cache Performance with Inclusive Caches: Temporal Locality Aware (TLA) Cache Management Policies .....	151
<i>Aamer Jaleel, Eric Borch, Malini Bhandaru, Simon C. Steely Jr., and Joel Emer</i>	
STEM: Spatiotemporal Management of Capacity for Intra-core Last Level Caches .....	163
<i>Dongyuan Zhan, Hong Jiang, and Sharad C. Seth</i>	
Sampling Dead Block Prediction for Last-Level Caches .....	175
<i>Samira Manabi Khan, Yingying Tian, and Daniel A. Jiménez</i>	
The ZCache: Decoupling Ways and Associativity .....	187
<i>Daniel Sanchez and Christos Kozyrakis</i>	
<b>Session 3B: Data Parallelism</b>	
Efficient Selection of Vector Instructions Using Dynamic Programming .....	201
<i>Rajkishore Barik, Jisheng Zhao, and Vivek Sarkar</i>	
Many-Thread Aware Prefetching Mechanisms for GPGPU Applications .....	213
<i>Jaekyu Lee, Nagesh B. Lakshminarayana, Hyesoon Kim, and Richard Vuduc</i>	

Single-Chip Heterogeneous Computing: Does the Future Include Custom Logic, FPGAs, and GPGPUs? .....	225
<i>Eric S. Chung, Peter A. Milder, James C. Hoe, and Ken Mai</i>	

Improving SIMT Efficiency of Global Rendering Algorithms with Architectural Support for Dynamic Micro-Kernels .....	237
<i>Michael Steffen and Joseph Zambreno</i>	

## Session 4A: Concurrency

InstantCheck: Checking the Determinism of Parallel Programs Using On-the-Fly Incremental Hashing .....	251
<i>Adrian Nistor, Darko Marinov, and Josep Torrellas</i>	

Tolerating Concurrency Bugs Using Transactions as Lifeguards .....	263
<i>Jie Yu and Satish Narayanasamy</i>	

Architectural Support for Fair Reader-Writer Locking .....	275
<i>Enrique Vallejo, Ramón Beivide, Adrián Cristal, Tim Harris, Fernando Vallejo, Osman Unsal, and Mateo Valero</i>	

AtomTracker: A Comprehensive Approach to Atomic Region Inference and Violation Detection .....	287
<i>Abdullah Muzahid, Norimasa Otsuki, and Josep Torrellas</i>	

## Session 4B: Microarchitecture I

Register Cache System Not for Latency Reduction Purpose .....	301
<i>Ryota Shioya, Kazuo Horio, Masahiro Goshima, and Shuichi Sakai</i>	

Synergistic TLBs for High Performance Address Translation in Chip Multiprocessors .....	313
<i>Shekhar Srikantaiah and Mahmut Kandemir</i>	

Erasing Core Boundaries for Robust and Configurable Performance .....	325
<i>Shantanu Gupta, Shuguang Feng, Amin Ansari, and Scott Mahlke</i>	

Minimal Multi-threading: Finding and Removing Redundant Instructions in Multi-threaded Processors .....	337
<i>Guoping Long, Diana Franklin, Susmit Biswas, Pablo Ortiz, Jason Oberg, Dongrui Fan, and Frederic T. Chong</i>	

## Session 5A: Memories

Parichute: Generalized Turbocode-Based Error Correction for Near-Threshold Caches .....	351
<i>Timothy N. Miller, Renji Thomas, James Dinan, Bruce Adcock, and Radu Teodorescu</i>	

Understanding the Energy Consumption of Dynamic Random Access Memories .....	363
<i>Thomas Vogelsang</i>	

Elastic Refresh: Techniques to Mitigate Refresh Penalties in High Density Memory .....	375
<i>Jeffrey Stuecheli, Dimitris Kaseridis, Hillery C.Hunter, and Lizy K. John</i>	

Moneta: A High-Performance Storage Array Architecture for Next-Generation, Non-volatile Memories .....	385
<i>Adrian M. Caulfield, Arup De, Joel Coburn, Todor I. Mollow, Rajesh K. Gupta, and Steven Swanson</i>	

## Session 5B: NoCs

Pseudo-Circuit: Accelerating Communication for On-Chip Interconnection Networks .....	399
<i>Minseon Ahn and Eun Jung Kim</i>	
LOFT: A High Performance Network-on-Chip Providing Quality-of-Service Support .....	409
<i>Jin Ouyang and Yuan Xie</i>	
Throughput-Effective On-Chip Networks for Manycore Accelerators .....	421
<i>Ali Bakhoda, John Kim, and Tor M. Aamodt</i>	
Adaptive Flow Control for Robust Performance and Energy .....	433
<i>Syed Ali Raza Jafri, Yu-Ju Hong, Mithuna Thottethodi, and T.N. Vijaykumar</i>	

## Session 6A: Coherence

ScalableBulk: Scalable Cache Coherence for Atomic Blocks in a Lazy Environment .....	447
<i>Xuehai Qian, Wonsun Ahn, and Josep Torrellas</i>	
Virtual Snooping: Filtering Snoops in Virtualized Multi-cores .....	459
<i>Daehoon Kim, Hwanju Kim, and Jaehyuk Huh</i>	
Fractal Coherence: Scalably Verifiable Cache Coherence .....	471
<i>Meng Zhang, Alvin R. Lebeck, and Daniel J. Sorin</i>	

## Session 6B: Microarchitecture II

A Predictive Model for Dynamic Microarchitectural Adaptivity Control .....	485
<i>Christophe Dubach, Timothy M. Jones, Edwin V. Bonilla, and Michael F.P. O'Boyle</i>	
ReMAP: A Reconfigurable Heterogeneous Multicore Architecture .....	497
<i>Matthew A. Watkins and David H. Albonesi</i>	
Probabilistic Distance-Based Arbitration: Providing Equality of Service for Many-Core CMPs .....	509
<i>Michael M. Lee, John Kim, Dennis Abts, Michael Marty, and Jae W. Lee</i>	

## Session 7: Tools

Adaptive and Speculative Slack Simulations of CMPs on CMPs .....	523
<i>Jainwei Chen, Lakshmi Kumar Dabbiru, Daniel Wong, Murali Annavaram, and Michel Dubois</i>	
SD3: A Scalable Approach to Dynamic Data-Dependence Profiling .....	535
<i>Minjang Kim, Hyesoon Kim, and Chi-Keung Luk</i>	

Automatic Parallelization in a Binary Rewriter .....547  
*Aparna Kotha, Kapil Anand, Matthew Smithson, Greeshma Yellareddy, and Rajeev Barua*

**Author Index** .....559