

Proceedings of the

The 40th Annual IEEE/ACM International Symposium on Microarchitecture

December 1-5, 2007, Chicago, Illinois, USA



Los Alamitos, California

Washington • Tokyo



Table of Contents

40th IEEE/ACM International Symposium on Microarchitecture (MICRO 2007)

Message from the General Chairs	viii
Message from the Program Chairs	ix
Organizing Committee	x
Reviewers	xi

Session 1: Technology Issues

Optimizing NUCA Organizations and Wiring Alternatives for Large Caches with CACTI 6.0	3
<i>Naveen Muralimanohar, Rajeev Balasubramonian, and Norm Jouppi</i>	
Process Variation Tolerant 3T1D-Based Cache Architectures	15
<i>Xiaoyao Liang, Ramon Canal, Gu-Yeon Wei, and David Brooks</i>	
Mitigating Parameter Variation with Dynamic Fine-Grain Body Biasing	27
<i>Radu Teodorescu, Jun Nakano, Abhishek Tiwari, and Josep Torrellas</i>	

Session 2A: Instruction Scheduling

Optimal versus Heuristic Global Code Scheduling	43
<i>Sebastian Winkel</i>	
Global Multi-Threaded Instruction Scheduling	56
<i>Guilherme Ottoni and David August</i>	
Revisiting the Sequential Programming Model for Multi-Core	69
<i>Matthew Bridges, Neil Vachharajani, Yun Zhang, Thomas Jablin, and David August</i>	

Session 2B: Wear-Out Aware Architectures

Penelope: The NBTI-Aware Processor	85
<i>Jaume Abella, Xavier Vera, and Antonio González</i>	
Software-Based Online Detection of Hardware Defects: Mechanisms, Architectural Support, and Evaluation	97
<i>Kypros Constantinides, Onur Mutlu, Todd Austin, and Valeria Bertacco</i>	
Self-calibrating Online Wearout Detection	109
<i>Jason Blome, Shuguang Feng, Shantanu Gupta, and Scott Mahlke</i>	

Session 3A: Memory

Implementing Signatures for Transactional Memory	123
<i>Daniel Sanchez, Luke Yen, Mark D. Hill, and Karthikeyan Sankaralingam</i>	
Smart Refresh: An Enhanced Memory Controller Design for Reducing Energy in Conventional and 3D Die-Stacked DRAMs	134
<i>Mrinmoy Ghosh and Hsien-Hsin S. Lee</i>	

Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors _____	146
<i>Onur Mutlu and Thomas Moscibroda</i>	

Session 3B: Networking and Security

Impact of Cache Coherence Protocols on the Processing of Network Traffic _____	161
<i>Amit Kumar and Ram Huggahalli</i>	

Flattened Butterfly Topology for On-Chip Networks _____	172
<i>John Kim, James Balfour, and William Dally</i>	

Using Address Independent Seed Encryption and Bonsai Merkle Trees to Make Secure Processors OS- and Performance-Friendly _____	183
<i>Brian Rogers, Siddhartha Chhabra, Milos Prvulovic, and Yan Solihin</i>	

Session 4A: Reliability

Multi-bit Error Tolerant Caches Using Two-Dimensional Error Coding _____	197
<i>Jangwoo Kim, Nikos Hardavellas, Ken Mai, Babak Falsafi, and James Hoe</i>	

Argus: Low-Cost, Comprehensive Error Detection in Simple Cores _____	210
<i>Albert Meixner, Michael E. Bauer, and Daniel Sorin</i>	

Leveraging 3D Technology for Improved Reliability _____	223
<i>Niti Madan and Rajeev Balasubramonian</i>	

Effective Optimistic-Checker Tandem Core Design through Architectural Pruning _____	236
<i>Francisco Mesa-Martínez and Jose Renau</i>	

Session 4B: Simulation/Workload Analysis

FPGA-Accelerated Simulation Technologies (FAST): Fast, Full-System, Cycle-Accurate Simulators _____	249
<i>Derek Chiou, Dam Sunwoo, Joonsoo Kim, Nikhil A. Patil, William Reinhart, Darrel Eric Johnson, Jebediah Keefe, and Hari Angepat</i>	

Microarchitectural Design Space Exploration Using an Architecture-Centric Approach _____	262
<i>Christophe Dubach, Timothy Jones, and Michael O'Boyle</i>	

Informed Microarchitecture Design Space Exploration Using Workload Dynamics _____	274
<i>Chang-Burm Cho, Wangyuan Zhang, and Tao Li</i>	

Time Interpolation: So Many Metrics, So Few Registers _____	286
<i>Todd Mytkowicz, Peter F. Sweeney, Matthias Hauswirth, and Amer Diwan</i>	

Session 5: Prefetching and Snooping

Low-Cost Epoch-Based Correlation Prefetching for Commercial Applications _____	301
<i>Yuan Chou</i>	

A Framework for Coarse-Grain Optimizations in the On-Chip Memory Hierarchy _____	314
<i>Jason Zebchuk, Elham Safi, and Andreas Moshovos</i>	

Uncorq: Unconstrained Snoop Request Delivery in Embedded-Ring Multiprocessors _____	327
<i>Karin Strauss, Xiaowei Shen, and Josep Torrellas</i>	

Session 6: Parallelism and QoS in CMPs

A Framework for Providing Quality of Service in Chip Multi-Processors _____	343
<i>Fei Guo, Yan Solihin, Li Zhao, and Ravishankar Iyer</i>	
A Practical Approach to Exploiting Coarse-Grained Pipeline Parallelism in C Programs _____	356
<i>William Thies, Vikram Chandrasekhar, and Saman Amarasinghe</i>	
Data Access Partitioning for Fine-grain Parallelism on Multicore Architectures _____	369
<i>Michael Chu, Rajiv Ravindran, and Scott Mahlke</i>	

Session 7: Parallel Architectures

Composable Lightweight Processors _____	381
<i>Changkyu Kim, Simha Sethumadhavan, M.S. Govindan, Nitya Ranganathan, Divya Gulati, Doug Burger, and Stephen W. Keckler</i>	
The Art of Deception: Adaptive Precision Reduction for Area Efficient Physics Acceleration _____	394
<i>Thomas Yeh, Petros Faloutsos, Milos Ercegovac, Sanjay Patel, and Glenn Reinman</i>	
Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow _____	407
<i>Wilson W.L. Fung, Ivan Sham, George Yuan, and Tor M. Aamodt</i>	

Session 8: Cache Replacement Policies

Scavenger: A New Last Level Cache Architecture with Global Block Priority _____	421
<i>Arkaprava Basu, Nevin Kirman, Meyrem Kirman, Mainak Chaudhuri, and José Martínez</i>	
Guaranteeing Hits to Improve the Efficiency of a Small Instruction Cache _____	433
<i>Stephen Hines, David Whalley, and Gary Tyson</i>	
Emulating Optimal Replacement with a Shepherd Cache _____	445
<i>Kaushik Rajan and Govindarajan Ramaswamy</i>	

Author Index _____	455
---------------------------	------------