

# **2013 46th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2013)**

**Davis, California, USA  
7 – 11 December 2013**



**IEEE Catalog Number: CFP13071-POD  
ISBN: 978-1-5090-6603-2**

**Copyright © 2013, The 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:	CFP13071-POD
ISBN (Print-On-Demand):	978-1-5090-6603-2
ISBN (Online):	978-1-4503-2638-4
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](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# The 46<sup>th</sup> Annual IEEE/ACM International Symposium on Microarchitecture MICRO 2013

## Table of Contents

Message from the General Chair.....	vii
Message from the Program Chair.....	viii
Organizing Committee.....	x
Program Committee.....	xi
External Reviewers.....	xii

---

### Session 1A - Approximate Computing

Quality-Programmable Vector Processors for Approximate Computing.....	1
<i>Swagath Venkataramani, Vinay K. Chippa, Srimat T. Chakradhar, Kaushik Roy, and Anand Raghunathan</i>	
SAGE: Self-Tuning Approximation for Graphics Engines.....	13
<i>Mehrzad Samadi, Janghaeng Lee, Davoud Jamshidi, Amir Hormati, and Scott Mahlke</i>	
Approximate Storage in Solid-State Memories.....	25
<i>Adrian Sampson, Jacob Nelson, Karin Strauss, and Luis Ceze</i>	

### Session 1B - Energy Optimizations

MLP-Aware Dynamic Instruction Window Resizing for Adaptively Exploiting Both ILP and MLP.....	37
<i>Yuya Kora, Kyohei Yamaguchi, and Hideki Ando</i>	
TLC: A Tag-less Cache Design for Reducing Dynamic First Level Cache Energy .....	49
<i>Andreas Sembrant, Erik Hagersten, and David Black-Schaffer</i>	
Decoupled Compressed Cache: Exploiting Spatial Locality for Energy-Optimized Compressed Caching.....	62
<i>Somayeh Sardashti and David A. Wood</i>	

## Session 2A - Power-Efficient GPUs

Exploiting GPU Peak-power and Performance Tradeoffs through Reduced Effective Pipeline Latencies.....	74
<i>Syed Zohaib Gilani, Nam Sung Kim, and Michael Schulte</i>	
A Locality-Aware Memory Hierarchy for Energy-Efficient GPU Architectures.....	86
<i>Minsoo Rhu, Michael Sullivan, Jingwen Leng, and Mattan Erez</i>	
Divergence-Aware Warp Scheduling.....	99
<i>Timothy G. Rogers, Mike O'Connor, and Tor M. Aamodt</i>	
Warped Gates: Gating Aware Scheduling and Power Gating for GPGPUs.....	111
<i>Mohammad Abdel-Majeed, Daniel Wong, and Murali Annavaram</i>	

## Session 2B – Resilience

Virtually Aged Sampling DMR: Unifying Circuit Failure Detection and Circuit Failure Prediction.....	123
<i>Raghuraman Balasubramanian and Karthikeyan Sankaralingam</i>	
Use It Or Lose It: Wear-out and Lifetime in Future Chip Multiprocessors.....	136
<i>Hyungjun Kim, Arseniy Vitkovskiy, Paul V. Gratz, and Vassos Soteriou</i>	
uDIREC: Unified Diagnosis and Reconfiguration for Frugal Bypass of NoC faults .....	148
<i>Ritesh Parikh and Valeria Bertacco</i>	
Implicit Storing and Redundant Encoding of Attribute Information in Error Correction Codes.....	160
<i>Yiannakis Sazeides, Emre Ozer (ARM), Danny Kershaw, Panagiota Nikolaou, Marios Kleanthous, and Jaume Abella</i>	

## Session 3A – Main Memory

Linearly Compressed Pages: A Low-Complexity, Low-Latency Main Memory Compression Framework.....	172
<i>Gennady Pekhimenko, Vivek Seshadri, Yoongu Kim, Hongyi Xin, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry</i>	
RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization.....	185
<i>Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry</i>	
Quantifying the Relationship between the Power Delivery Network and Architectural Policies in a 3D-Stacked Memory Device.....	198
<i>Manjunath Shevgoor, Jung-Sik Kim, Niladrish Chatterjee, Rajeev Balasubramonian, Al Davis, and Aniruddha N. Udipi</i>	

### **Session 3B – Power Management & Interconnects**

Crank It Up or Dial It Down: Coordinated Multiprocessor Frequency and Folding Control.....210  
*Augusto Vega, Alper Buyuktosunoglu, Heather Hanson, Pradip Bose, and Srinivasan Ramani*

Wavelength Stealing: An Opportunistic Approach to Channel Sharing in  
Multi-chip Photonic Interconnects.....222  
*Arslan Zulfiqar, Pranay Koka, Herb Schwetman, Mikko Lipasti,  
Xuezhe Zheng, and Ashok V. Krishnamoorthy*

DESC: Energy-Efficient Data Exchange using Synchronized Counters.....234  
*Mahdi Nazm Bojnordi and Engin Ipek*

### **Session 4A – Prefetching**

Linearizing Irregular Memory Accesses for Improved Correlated Prefetching.....247  
*Akanksha Jain and Calvin Lin*

RDIP: Return-address-stack Directed Instruction Prefetching.....260  
*Aasheesh Kolli, Ali Saidi, and Thomas Wenisch*

SHIFT: Shared History Instruction Fetch for Lean-Core Server Processors.....272  
*Cansu Kaynak, Boris Grot, and Babak Falsafi*

### **Session 4B – Cache Management**

Insertion and Promotion for Tree-Based PseudoLRU Last-Level Caches.....284  
*Daniel A. Jimenez (Texas A&M)*

Imbalanced Cache Partitioning for Balanced Data-Parallel Programs.....297  
*Abhisek Pan and Vijay S. Pai*

The Reuse Cache: Downsizing the Shared Last-Level Cache.....310  
*Jorge Albericio, Pablo Ibáñez, Víctor Viñals, and José M. Llberia*

### **Session 5A – Programming, Compilation, and Provisioning**

Enabling Datacenter Servers to Scale Out Economically and Sustainably.....322  
*Chao Li, Yang Hu, Ruijin Zhou, Ming Liu, Longjun Liu, Jingling Yuan, and Tao Li*

Efficient Multiprogramming for Multicores with SCAF.....334  
*Timothy Creech, Aparna Kotha, and Rajeev Barua*

Allocating Rotating Registers by Scheduling.....346  
*Hongbo Rong, Hyunchul Park, Cheng Wang, and Youfeng Wu*

## **Session 5B – Coherence & Memory Management**

Multi-Grain Coherence Directory.....	359
<i>Jason Zebchuk, Babak Falsafi, and Andreas Moshovos</i>	
BulkCommit: Scalable and Fast Commit of Atomic Blocks in a Lazy Multiprocessor Environment.....	371
<i>Xuehai Qian, Benjamin Sahelics, Josep Torrellas, and Depei Qian</i>	
Large-Reach Memory Management Unit Caches.....	383
<i>Abhishek Bhattacharjee</i>	

## **Session 6A – GPU Memory Management**

Efficient Management of Last-level Caches in Graphics Processors for 3D Scene Rendering Workloads.....	395
<i>Jayesh Gaur, Raghuram Onti Srinivasan, Sreenivas Subramoney, and Mainak Chaudhuri</i>	
Energy Efficient GPU Transactional Memory via Space-Time Optimizations.....	408
<i>Wilson Wai Lun Fung and Tor M. Aamodt</i>	

## **Session 6B – Storage Optimizations**

Kiln: Closing the Performance Gap Between Systems With and Without Persistence Support.....	421
<i>Jishen Zhao, Sheng Li, Doe Hyun Yoon, Yuan Xie, and Norman P. Jouppi</i>	
Aegis: Partitioning Data Block for Efficient Recovery of Stuck-At-Faults in Phase Change Memory.....	433
<i>Jie Fan, Song Jiang, Jiwu Shu, Youhui Zhang, and Weimin Zhen</i>	

## **Session 7 – Heterogeneous Computing**

Trace Based Phase Prediction For Tightly-Coupled Heterogeneous Cores.....	445
<i>Shruti Padmanabha, Andrew Lukefahr, Reetuparna Das, and Scott Mahlke</i>	
Heterogeneous System Coherence for Integrated CPU-GPU Systems.....	457
<i>Jason Power, Arkaprava Basu, Junli Gu, Sooraj Puthoor, Bradford M. Beckmann, Mark D. Hill Steven K. Reinhardt, and David A. Wood</i>	
Meet the Walkers: Accelerating Index Traversals for In-Memory Databases.....	468
<i>Onur Kocberber, Boris Grot, Javier Picorel, Babak Falsafi, Kevin Lim, and Parthasarathy Ranganathan</i>	

<b>Author Index</b> .....	480
---------------------------	-----