

# **2017 IEEE International Symposium on High Performance Computer Architecture (HPCA 2017)**

**Austin, Texas, USA  
4 – 8 February 2017**



**IEEE Catalog Number: CFP17013-POD**  
**ISBN: 978-1-5090-4986-8**

**Copyright © 2017 by the Institute of Electrical and Electronics Engineers, Inc  
All Rights Reserved**

*Copyright and Reprint Permissions:* Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. 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:	CFP17013-POD
ISBN (Print-On-Demand):	978-1-5090-4986-8
ISBN (Online):	978-1-5090-4985-1
ISSN:	1530-0897

**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

# 2017 IEEE International Symposium on High Performance Computer Architecture

## HPCA 2017

### Table of Contents

Message from the Program Chair.....	xi
Committees.....	xiii
Program Committee.....	xiv
External Review Committee.....	xvi
Sponsors.....	xix

---

#### **HPCA Session 1: Lightning Rounds**

#### **HPCA Session 2: Best Paper Nominees**

Towards Pervasive and User Satisfactory CNN across GPU Microarchitectures .....	1
<i>Mingcong Song, Yang Hu, Huixiang Chen, and Tao Li</i>	
Near-Optimal Access Partitioning for Memory Hierarchies with Multiple Heterogeneous Bandwidth Sources .....	13
<i>Jayesh Gaur, Mainak Chaudhuri, Pradeep Ramachandran, and Sreenivas Subramoney</i>	
NCAP: Network-Driven, Packet Context-Aware Power Management for Client-Server Architecture .....	25
<i>Mohammad Alian, Ahmed H. M. O. Abulila, Lokesh Jindal, Daehoon Kim, and Nam Sung Kim</i>	
Supporting Address Translation for Accelerator-Centric Architectures .....	37
<i>Yuchen Hao, Zhenman Fang, Glenn Reinman, and Jason Cong</i>	

## **HPCA Session 3A: Industrial Session**

Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques .....	49
<i>Yu Cai, Saugata Ghose, Yixin Luo, Ken Mai, Onur Mutlu, and Erich F. Haratsch</i>	
Defect Analysis and Cost-Effective Resilience Architecture for Future DRAM Devices .....	61
<i>Sanguhn Cha, Seongil O, Hyunsung Shin, Sangjoon Hwang, Kwangil Park, Seong Jin Jang, Joo Sun Choi, Gyo Young Jin, Young Hoon Son, Hyunyeon Cho, Jung Ho Ahn, and Nam Sung Kim</i>	
Architecting an Energy-Efficient DRAM System for GPUs .....	73
<i>Niladrish Chatterjee, Mike O'Connor, Donghyuk Lee, Daniel R. Johnson, Stephen W. Keckler, Minsoo Rhu, and William J. Dally</i>	
Design and Analysis of an APU for Exascale Computing .....	85
<i>Thiruvengadam Vijayaraghavany, Yasuko Eckert, Gabriel H. Loh, Michael J. Schulte, Mike Ignatowski, Bradford M. Beckmann, William C. Brantley, Joseph L. Greathouse, Wei Huang, Arun Karunanithi, Onur Kayiran, Mitesh Meswani, Indrani Paul, Matthew Poremba, Steven Raasch, Steven K. Reinhardt, Greg Sadowski, and Vilas Sridharan</i>	
BRAVO: Balanced Reliability-Aware Voltage Optimization .....	97
<i>Karthik Swaminathan, Nandhini Chandramoorthy, Chen-Yong Cher, Ramon Bertran, Alper Buyuktosunoglu, and Pradip Bose</i>	

## **HPCA Session 3B: Cache**

Maximizing Cache Performance Under Uncertainty .....	109
<i>Nathan Beckmann and Daniel Sanchez</i>	
SWAP: Effective Fine-Grain Management of Shared Last-Level Caches with Minimum Hardware Support .....	121
<i>Xiaodong Wang, Shuang Chen, Jeff Setter, and José F. Martínez</i>	
A Split Cache Hierarchy for Enabling Data-Oriented Optimizations .....	133
<i>Andreas Sembrant, Erik Hagersten, and David Black-Schaffer</i>	
Fast and Accurate Exploration of Multi-level Caches Using Hierarchical Reuse Distance .....	145
<i>Rafael K. V. Maeda, Qiong Cai, Jiang Xu, Zhe Wang, and Zhongyuan Tian</i>	

## **HPCA Session 4A: Power, Energy & Large-Scale Computing**

Enabling Effective Module-Oblivious Power Gating for Embedded Processors .....	157
<i>Hari Cherupalli, Henry Duwe, Weidong Ye, Rakesh Kumar, and John Sartori</i>	

Application-Specific Performance-Aware Energy Optimization on Android Mobile Devices .....	169
<i>Karthik Rao, Jun Wang, Sudhakar Yalamanchili, Yorai Wardi, and Ye Handong</i>	
Fast Decentralized Power Capping for Server Clusters .....	181
<i>Reza Azimi, Masoud Badiei, Xin Zhan, Na Li, and Sherief Reda</i>	
Random Folded Clos Topologies for Datacenter Networks .....	193
<i>Cristóbal Camarero, Carmen Martínez, and Ramón Beivide</i>	

## **HPCA Session 4B: Memory**

Tiny Directory: Efficient Shared Memory in Many-Core Systems with Ultra-Low-Overhead Coherence Tracking .....	205
<i>Sudhanshu Shukla and Mainak Chaudhuri</i>	
Partial Row Activation for Low-Power DRAM System .....	217
<i>Yebin Lee, Hyeonggyu Kim, Seokin Hong, and Soontae Kim</i>	
Understanding and Optimizing Power Consumption in Memory Networks .....	229
<i>Xun Jian, Pavan Kumar Hanumolu, and Rakesh Kumar</i>	
SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies .....	241
<i>Hasan Hassan, Nandita Vijaykumar, Samira Khan, Saugata Ghose, Kevin Chang, Gennady Pekhimenko, Donghyuk Lee, Oguz Ergin, and Onur Mutlu</i>	

## **HPCA Session 5A: NOC**

Static Bubble: A Framework for Deadlock-Free Irregular On-chip Topologies .....	253
<i>Aniruddh Ramrakhyani and Tushar Krishna</i>	
Designing Low-Power, Low-Latency Networks-on-Chip by Optimally Combining Electrical and Optical Links .....	265
<i>Sebastian Werner, Javier Navaridas, and Mikel Luján</i>	
Near-Ideal Networks-on-Chip for Servers .....	277
<i>Pejman Lotfi-Kamran, Mehdi Modarressi, and Hamid Sarbazi-Azad</i>	
Design and Evaluation of AWGR-Based Photonic NoC Architectures for 2.5D Integrated High Performance Computing Systems .....	289
<i>Paolo Grani, Roberto Proietti, Venkatesh Akella, and S. J. Ben Yoo</i>	

## **HPCA Session 5B: Security**

Secure Dynamic Memory Scheduling Against Timing Channel Attacks .....	301
<i>Yao Wang, Benjamin Wu, and G. Edward Suh</i>	

Cold Boot Attacks are Still Hot: Security Analysis of Memory Scramblers in Modern Processors .....	313
<i>Salessawi Ferede Yitbarek, Misiker Tadesse Aga, Reetuparna Das, and Todd Austin</i>	

Cooperative Path-ORAM for Effective Memory Bandwidth Sharing in Server Settings .....	325
<i>Rujia Wang, Youtao Zhang, and Jun Yang</i>	

Camouflage: Memory Traffic Shaping to Mitigate Timing Attacks .....	337
<i>Yanqi Zhou, Sameer Wagh, Prateek Mittal, and David Wentzclaff</i>	

## **HPCA Session 6A: Emerging Storage**

SILC-FM: Subblocked InterLeaved Cache-Like Flat Memory Organization .....	349
<i>Jee Ho Ryoo, Mitesh R. Meswani, Andreas Prodromou, and Lizy K. John</i>	

ATOM: Atomic Durability in Non-volatile Memory through Hardware Logging .....	361
<i>Arpit Joshi, Vijay Nagarajan, Stratis Viglas, and Marcelo Cintra</i>	

KAML: A Flexible, High-Performance Key-Value SSD .....	373
<i>Yanqin Jin, Hung-Wei Tseng, Yannis Papakonstantinou, and Steven Swanson</i>	

Balancing Performance and Lifetime of MLC PCM by Using a Region Retention Monitor .....	385
<i>Mingzhe Zhang, Lunkai Zhang, Lei Jiang, Zhiyong Liu, and Frederic T. Chong</i>	

## **HPCA Session 6B: Scheduling**

Reliability-Aware Scheduling on Heterogeneous Multicore Processors .....	397
<i>Ajeya Naithani, Stijn Eyerman, and Lieven Eeckhout</i>	

Hipster: Hybrid Task Manager for Latency-Critical Cloud Workloads .....	409
<i>Rajiv Nishtala, Paul Carpenter, Vinicius Petrucci, and Xavier Martorell</i>	

Cooper: Task Colocation with Cooperative Games .....	421
<i>Qiuyun Llull, Songchun Fan, Seyed Majid Zahedi, and Benjamin C. Lee</i>	

MemPod: A Clustered Architecture for Efficient and Scalable Migration in Flat Address Space Multi-level Memories .....	433
<i>Andreas Prodromou, Mitesh Meswani, Nuwan Jayasena, Gabriel Loh, and Dean M. Tullsen</i>	

## **HPCA Session 7A: Novel Architectures**

Exploring Hyperdimensional Associative Memory .....	445
<i>Mohsen Imani, Abbas Rahimi, Deqian Kong, Tajana Rosing, and Jan M. Rabaey</i>	

GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks .....	457
<i>Lifeng Nai, Ramyad Hadidi, Jaewoong Sim, Hyojong Kim, Pranith Kumar, and Hyesoon Kim</i>	
High-Bandwidth Low-Latency Approximate Interconnection Networks .....	469
<i>Daichi Fujiki, Kiyo Ishii, Ikki Fujiwara, Hiroki Matsutani, Hideharu Amano, Henri Casanova, and Michihiro Koibuchi</i>	
Compute Caches .....	481
<i>Shaizeen Aga, Supreet Jeloka, Arun Subramaniyan, Satish Narayanasamy, David Blaauw, and Reetuparna Das</i>	
<b>HPCA Session 7B: Control-Flow and Microarchitecture</b>	
Boomerang: A Metadata-Free Architecture for Control Flow Delivery .....	493
<i>Rakesh Kumar, Cheng-Chieh Huang, Boris Grot, and Vijay Nagarajan</i>	
PABST: Proportionally Allocated Bandwidth at the Source and Target .....	505
<i>Derek R. Hower, Harold W. Cain, and Carl A. Waldspurger</i>	
SOUP-N-SALAD: Allocation-Oblivious Access Latency Reduction with Asymmetric DRAM Microarchitectures .....	517
<i>Yuhwan Ro, Hyunyoong Cho, Eojin Lee, Daejin Jung, Young Hoon Son, Jung Ho Ahn, and Jae W. Lee</i>	
Transparent and Efficient CFI Enforcement with Intel Processor Trace .....	529
<i>Yutao Liu, Peitao Shi, Xinran Wang, Haibo Chen, Binyu Zang, and Haibing Guan</i>	
<b>HPCA Session 8A: Accelerators</b>	
PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning .....	541
<i>Linghao Song, Xuehai Qian, Hai Li, and Yiran Chen</i>	
FlexFlow: A Flexible Dataflow Accelerator Architecture for Convolutional Neural Networks .....	553
<i>Wenyan Lu, Guihai Yan, Jiajun Li, Shijun Gong, Yinhe Han, and Xiaowei Li</i>	
Needle: Leveraging Program Analysis to Analyze and Extract Accelerators from Whole Programs .....	565
<i>Snehasish Kumar, Nick Sumner, Vijayalakshmi Srinivasan, Steve Margerm, and Arrvindh Shriraman</i>	
Radiation-Induced Error Criticality in Modern HPC Parallel Accelerators .....	577
<i>Daniel Alfonso Goncalves De Oliveira, Laercio Lima Pilla, Mauricio Hanzich, Vinicius Fratin, Fernando Fernandes, Caio Lunardi, José María Cela, Philippe Olivier Alexandre Navaux, Luigi Carro, and Paolo Rech</i>	

## **HPCA Session 8B: Best of CAL**

### **HPCA Session 9A: GPU Power & Energy**

Pilot Register File: Energy Efficient Partitioned Register File for GPUs .....	589
<i>Mohammad Abdel-Majeed, Alireza Shafaei, Hyeran Jeon, Massoud Pedram, and Murali Annavaram</i>	
G-Scalar: Cost-Effective Generalized Scalar Execution Architecture for Power-Efficient GPUs .....	601
<i>Zhenhong Liu, Syed Gilani, Murali Annavaram, and Nam Sung Kim</i>	
Dynamic GPGPU Power Management Using Adaptive Model Predictive Control .....	613
<i>Abhinandan Majumdar, Leonardo Piga, Indrani Paul, Joseph L. Greathouse, Wei Huang, and David H. Albonesi</i>	

### **HPCA Session 9B: GPU**

Efficient Sequential Consistency in GPUs via Relativistic Cache Coherence .....	625
<i>Xiaowei Ren and Mieszko Lis</i>	
Processing-in-Memory Enabled Graphics Processors for 3D Rendering .....	637
<i>Chenhao Xie, Shuaiwen Leon Song, Jing Wang, Weigong Zhang, and Xin Fu</i>	
Controlled Kernel Launch for Dynamic Parallelism in GPUs .....	649
<i>Xulong Tang, Ashutosh Pattnaik, Huaipan Jiang, Onur Kayiran, Adwait Jog, Sreepathi Pai, Mohamed Ibrahim, Mahmut T. Kandemir, and Chita R. Das</i>	

<b>Author Index</b> .....	661
---------------------------	-----