

2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2018)

**Fukuoka, Japan
20-24 October 2018**

Pages 1-493



**IEEE Catalog Number: CFP18071-POD
ISBN: 978-1-5386-6241-0**

**Copyright © 2018 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:	CFP18071-POD
ISBN (Print-On-Demand):	978-1-5386-6241-0
ISBN (Online):	978-1-5386-6240-3

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

2018 51st Annual IEEE/ACM International Symposium on Microarchitecture **MICRO 2018**

Table of Contents

Message from the General Chairs	xvi
Message from the Program Chairs	xvii
Organization Committee	xx
Program Committee	xxi
Reviewers	xxiv

1A: Accelerators

Exploiting Locality in Graph Analytics through Hardware-Accelerated Traversal Scheduling	1
<i>Anurag Mukkara (MIT CSAIL), Nathan Beckmann (CMU SCS), Maleen Abeydeera (MIT CSAIL), Xiaosong Ma (QCRI), and Daniel Sanchez (MIT CSAIL)</i>	
Cambricon-S: Addressing Irregularity in Sparse Neural Networks through A Cooperative Software/Hardware Approach	15
<i>Xuda Zhou (University of Science and Technology of China, China), Zidong Du (Intelligence Processor Research Center, Institute of Computing Technology, CAS, China), Qi Guo (Intelligence Processor Research Center, Institute of Computing Technology, CAS, China), Shaoli Liu (Institute of Computing Technology, CAS), Chengsi Liu (Michigan State University, America), Chao Wang (University of Science and Technology of China, China), Xuehai Zhou (University of Science and Technology of China, China), Ling Li (Institute of Software, CAS, China), Tianshi Chen (Intelligence Processor Research Center, Institute of Computing Technology, CAS, China), and Yunji Chen (Intelligence Processor Research Center, Institute of Computing Technology, CAS, China)</i>	
CSE: Parallel Finite State Machines with Convergence Set Enumeration	29
<i>Youwei Zhuo (University of Southern California), Jinglei Cheng (University of Southern California), Qinyi Luo (University of Southern California), Jidong Zhai (Tsinghua University, China), Yanzhi Wang (Syracuse University), Zhongzhi Luan (Beihang University, China), and Xuehai Qian (University of Southern California)</i>	
Inter-Thread Communication in Multithreaded, Reconfigurable Coarse-Grain Arrays	42
<i>Dani Voitsechov (Technion Israel institute of technology, Israel), Oron Port (Technion Israel institute of technology, Israel), and Yoav Etsion@tce.technion.ac.il</i>	

An Architectural Framework for Accelerating Dynamic Parallel Algorithms on Reconfigurable Hardware .55.....	
	<i>Tao Chen (Cornell University), Shreesha Srinath (Cornell University), Christopher Batten (Cornell University), and G. Edward Suh (Cornell University)</i>

1B: Microarchitecture

Composable Building Blocks to Open up Processor Design .68.....	
	<i>Sizhuo Zhang (Massachusetts Institute of Technology), Andrew Wright (Massachusetts Institute of Technology), Thomas Bourgeat (Massachusetts Institute of Technology), and Arvind Arvind (Massachusetts Institute of Technology)</i>
Performance Improvement by Prioritizing the Issue of the Instructions in Unconfident Branch Slices .82.....	
	<i>Hideki Ando (Nagoya University, Japan)</i>
The Superfluous Load Queue .95.....	
	<i>Alberto Ros (Department of Computer Engineering, University of Murcia, Spain) and Stefanos Kaxiras (Department of Information Technology, Uppsala University, Sweden)</i>
Architectural Support for Probabilistic Branches .108.....	
	<i>Almutaz Adileh (Ghent University), David J. Lilja (University of Minnesota), and Lieven Eeckhout (Ghent University)</i>
STRAIGHT: Hazardless Processor Architecture Without Register Renaming .121.....	
	<i>Hidetsugu Irie (The University of Tokyo, Japan), Toru Koizumi (The University of Tokyo, Japan), Akifumi Fukuda (The University of Tokyo, Japan), Seiya Akaki (The University of Tokyo, Japan), Satoshi Nakae (The University of Tokyo, Japan), Yutaro Bessho (The University of Tokyo, Japan), Ryota Shioya (The University of Tokyo, Japan), Takahiro Notsu (FUJITSU LABORATORIES LTD., Japan), Katsuhiko Yoda (FUJITSU LABORATORIES LTD., Japan), Teruo Ishihara (FUJITSU LABORATORIES LTD., Japan), and Shuichi Sakai (The University of Tokyo, Japan)</i>

2 A: ML Accelerators

Diffy: a Déjà vu-Free Differential Deep Neural Network Accelerator .134.....	
	<i>Mostafa Mahmoud (University of Toronto, Canada), Kevin Siu (University of Toronto, Canada), and Andreas Moshovos (University of Toronto, Canada)</i>
Beyond the Memory Wall: A Case for Memory-Centric HPC System for Deep Learning .148.....	
	<i>Youngeun Kwon (KAIST) and Minsoo Rhu (KAIST)</i>
Towards Memory Friendly Long-Short Term Memory Networks (LSTMs) on Mobile GPUs .162.....	
	<i>Xingyao Zhang (University of Houston), Chenhao Xie (University of Houston), Jing Wang (Capital Normal University), Weidong Zhang (Capital Normal University), and Xin Fu (University of Houston)</i>

A Network-Centric Hardware/Algorithm Co-Design to Accelerate Distributed Training of Deep Neural Networks .175.....	
	<i>Youjie Li (University of Illinois at Urbana–Champaign, United States), Jongse Park (Georgia Institute of Technology, United States), Mohammad Alian (University of Illinois at Urbana–Champaign, United States), Yifan Yuan (University of Illinois at Urbana–Champaign, United States), Zheng Qu (Tsinghua University, China), Peitian Pan (Shanghai Jiao Tong University, China), Ren Wang (Intel Corporation, United States), Alexander Schwing (University of Illinois at Urbana–Champaign, United States), Hadi Esmaeilzadeh (University of California San Diego, United States), and Nam Sung Kim (University of Illinois at Urbana–Champaign, United States)</i>
PermDNN: Efficient Compressed DNN Architecture with Permuted Diagonal Matrices .189.....	
	<i>Chunhua Deng (City University of New York), Siyu Liao (City University of New York), Yi Xie (City University of New York), Keshab K. Parhi (University of Minnesota, Twin Cities), Xuehai Qian (University of Southern California), and Bo Yuan (City University of New York)</i>

2B: Compilers and Programming Languages

Rethinking the Memory Hierarchy for Modern Languages .203.....	
	<i>Po-An Tsai (Massachusetts Institute of Technology), Yee Ling Gan (Massachusetts Institute of Technology), and Daniel Sanchez (Massachusetts Institute of Technology)</i>
Harmonizing Speculative and Non-Speculative Execution in Architectures for Ordered Parallelism .217.....	
	<i>Mark C. Jeffrey (Massachusetts Institute of Technology), Victor A. Ying (Massachusetts Institute of Technology), Suvinay Subramanian (Massachusetts Institute of Technology), Hyun Ryong Lee (Massachusetts Institute of Technology), Joel Emer (NVIDIA/MIT), and Daniel Sanchez (Massachusetts Institute of Technology)</i>
Sampler: PMU-Based Sampling to Detect Memory Errors Latent in Production Software .231.....	
	<i>Sam Silvestro (University of Texas at San Antonio), Hongyu Liu (University of Texas at San Antonio), Tong Zhang (Virginia Tech), Changhee Jung (Virginia Tech), Dongyoon Lee (Virginia Tech), and Tongping Liu (University of Texas at San Antonio)</i>
TAPAS: Generating Parallel Accelerators from Parallel Programs .245.....	
	<i>Steven Margerm (Simon Fraser University), Amirali Sharifian (Simon Fraser University), Apala Guha (Simon Fraser University), Arrvindh Shriraman (Simon Fraser University), and Gilles Pokam (Intel Corporation)</i>
iDO: Compiler-Directed Failure Atomicity for Nonvolatile Memory .258.....	
	<i>Qingrui Liu (Virginia Tech), Joseph Izraelevitz (University of Rochester), Se Kwon Lee (UNIST), Michael L. Scott (University of Rochester), Sam H. Noh (UNIST), and Changhee Jung (Virginia Tech)</i>

3A: Memory Systems - I

Scalable Distributed Last-Level TLBs Using Low-Latency Interconnects .271.....	
<i>Srikant Bharadwaj (Georgia Institute of Technology and AMD Research), Guilherme Cox (Rutgers University), Tushar Krishna (Georgia Institute of Technology), and Abhishek Bhattacharjee (Rutgers University)</i>	
Duplicon Cache: Mitigating Off-Chip Memory Bank and Bank Group Conflicts Via Data Duplication .285.....	
<i>Ben Lin (University of Texas at Austin), Michael B. Healy (IBM T.J. Watson Research Center), Rustam Miftakhutdinov (Intel Corporation), Philip G. Emma (IBM T.J. Watson Research Center), and Yale Patt (University of Texas at Austin)</i>	
Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration .298.....	
<i>Yaohua Wang (ETH Zürich, Switzerland; National University of Defense Technology, China), Arash Tavakkol (ETH Zürich, Switzerland), Lois Orosa (ETH Zürich, Switzerland; University of Campinas, Brazil), Saugata Ghose (Carnegie Mellon University, USA), Nika Mansouri Ghiasi (ETH Zürich, Switzerland), Minesh Patel (ETH Zürich, Switzerland), Jeremie S. Kim (Carnegie Mellon University, USA; ETH Zürich, Switzerland), Hasan Hassan (ETH Zürich, Switzerland), Mohammad Sadrosadati (ETH Zürich, Switzerland), and Onur Mutlu (ETH Zürich, Switzerland; Carnegie Mellon University, USA)</i>	
CABLE: A CAche-Based Link Encoder for Bandwidth-Starved Manycores .312.....	
<i>Tri Nguyen (Princeton University, USA), Adi Fuchs (Princeton University, USA), and David Wentzlaff (Princeton University, USA)</i>	
Attaché: Towards Ideal Memory Compression by Mitigating Metadata Bandwidth Overheads .326....	
<i>Seokin Hong (IBM Thomas J. Watson Research Center), Prashant Jayaprakash Nair (IBM Thomas J. Watson Research Center), Bulent Abali (IBM Thomas J. Watson Research Center), Alper Buyuktosunoglu (IBM Thomas J. Watson Research Center), Kyu-Hyoun Kim (IBM Thomas J. Watson Research Center), and Michael Healy (IBM Thomas J. Watson Research Center)</i>	

3B: GPGPU/GPU

Combining HW/SW Mechanisms to Improve NUMA Performance of Multi-GPU Systems .339.....	
<i>Vinson Young (Georgia Institute of Technology), Aamer Jaleel (NVIDIA), Evgeny Bolotin (NVIDIA), Eiman Ebrahimi (NVIDIA), David Nellans (NVIDIA), and Oreste Villa (NVIDIA)</i>	
Neighborhood-Aware Address Translation for Irregular GPU Applications .352.....	
<i>Seunghye Shin (Binghamton University), Michael LeBeane (Advanced Micro Devices, Inc.), Yan Solihin (University of Central Florida), and Arkaprava Basu (Indian Institute of Science)</i>	
FineReg: Fine-Grained Register File Management for Augmenting GPU Throughput .364.....	
<i>Yunho Oh (Yonsei University, Korea), Myung Kuk Yoon (Yonsei University, Korea), William J. Song (Yonsei University, Korea), and Won Woo Ro (Yonsei University, Korea)</i>	

In-Register Parameter Caching for Dynamic Neural Nets with Virtual Persistent Processor Specialization .377.....	
<i>Farzad Khorasani (Tesla, Inc.), Hodjat Asghari Esfeden (University of California Riverside), Nael Abu-Ghazaleh (University of California Riverside), and Vivek Sarkar (Georgia Institute of Technology)</i>	
Voltage-Stacked GPUs: A Control Theory Driven Cross-Layer Solution for Practical Voltage Stacking in GPUs .390.....	
<i>An Zou (Washington University in St. Louis), Jingwen Leng (Shanghai Jiao Tong University), Xin He (Washington University in St. Louis), Yazhou Zu (University of Texas at Austin), Christopher D. Gill (Washington University in St. Louis), Vijay Janapa Reddi (Harvard University, The University of Texas at Austin), and Xuan Zhang (Washington University in St. Louis)</i>	

4A: Security - I

Osiris: A Low-Cost Mechanism to Enable Restoration of Secure Non-Volatile Memories .403.....	
<i>Mao Ye (University of Central Florida), Clayton Hughes (Sandia National Laboratories), and Amro Awad (University of Central Florida)</i>	
Morphable Counters: Enabling Compact Integrity Trees For Low-Overhead Secure Memories .416...	
<i>Gururaj Saileshwar (Georgia Institute of Technology, USA), Prashant Nair (IBM Research, USA), Prakash Ramrakhyani (Arm Research, USA), Wendy Elsasser (Arm Research, USA), Jose Joao (Arm Research, USA), and Moinuddin Qureshi (Georgia Institute of Technology, USA)</i>	
InvisiSpec: Making Speculative Execution Invisible in the Cache Hierarchy .428.....	
<i>Mengjia Yan (University of Illinois at Urbana-Champaign, USA), Jiho Choi (University of Illinois at Urbana-Champaign, USA), Dimitrios Skarlatos (University of Illinois at Urbana-Champaign, USA), Adam Morrison (Tel Aviv University, Israel), Christopher Fletcher (University of Illinois at Urbana-Champaign, USA), and Josep Torrellas (University of Illinois at Urbana-Champaign, USA)</i>	
Improving the Performance and Endurance of Encrypted Non-Volatile Main Memory through Deduplicating Writes .442.....	
<i>Pengfei Zuo (Huazhong University of Science and Technology, China), Yu Hua (Huazhong University of Science and Technology, China), Ming Zhao (Arizona State University, USA), Wen Zhou (Huazhong University of Science and Technology, China), and Yuncheng Guo (Huazhong University of Science and Technology, China)</i>	

4B: Storage Systems & Technologies

SSDcheck: Timely and Accurate Prediction of Irregular Behaviors in Black-Box SSDs .455.....	
<i>Joonsung Kim (Seoul National University, South Korea), Pyeongsu Park (Seoul National University, South Korea), Jaehyung Ahn (POSTECH, South Korea), Jihun Kim (POSTECH, South Korea), Jong Kim (POSTECH, South Korea), and Jangwoo Kim (Seoul National University, South Korea)</i>	

Amber*: Enabling Precise Full-System Simulation with Detailed Modeling of All SSD Resources .469.....	
<i>Donghyun Gouk (Yonsei University, Republic of Korea), Miryeong Kwon (Yonsei University, Republic of Korea), Jie Zhang (Yonsei University, Republic of Korea), Sungjoon Koh (Yonsei University, Republic of Korea), Wonil Choi (Yonsei University and Pennsylvania State University), Nam Sung Kim (University of Illinois Urbana-Champaign, USA), Mahmut Kandemir (Pennsylvania State University, USA), and Myoungsoo Jung (Yonsei University, Republic of Korea)</i>	
Invalid Data-Aware Coding to Enhance the Read Performance of High-Density Flash Memories .482	
<i>Wonil Choi (Pennsylvania State University, United States), Myoungsoo Jung (Yonsei University, Korea), and Mahmut Kandemir (Pennsylvania State University, United States)</i>	
Persistence Parallelism Optimization: A Holistic Approach from Memory Bus to RDMA Network .494.	
<i>Xing Hu (University of California, Santa Barbara, United States), Matheus Ogleari (University of California, Santa Cruz, United States), Jishen Zhao (University of California, San Diego, United States), Shuangchen Li (University of California, Santa Barbara, United States), Abanti Basak (University of California, Santa Barbara, United States), and Yuan Xie (University of California, Santa Barbara, United States)</i>	

5A: Memory Systems - II

PiCL: A Software-Transparent, Persistent Cache Log for Nonvolatile Main Memory .507.....	
<i>Tri Nguyen (Princeton University, USA) and David Wentzlaff (Princeton University, USA)</i>	
Efficient Hardware-Assisted Logging with Asynchronous and Direct-Update for Persistent Memory .520.....	
<i>Jungi Jeong (KAIST, South Korea), Chang Hyun Park (KAIST, South Korea), Jaehyuk Huh (KAIST, South Korea), and Seungryoul Maeng (KAIST, South Korea)</i>	
CHAMELEON: A Dynamically Reconfigurable Heterogeneous Memory System .533.....	
<i>Jagadish B. Kotra (AMD Research), Haibo Zhang (The Pennsylvania State University), Alaa R. Alameldeen (Intel Labs), Chris Wilkerson (NVIDIA Corporation Ltd), and Mahmut T. Kandemir (The Pennsylvania State University)</i>	
Compresso: Pragmatic Main Memory Compression .546.....	
<i>Esha Choukse (University of Texas at Austin), Mattan Erez (University of Texas at Austin), and Alaa R. Alameldeen (Intel Labs)</i>	
Farewell My Shared LLC! A Case for Private Die-Stacked DRAM Caches for Servers .559.....	
<i>Amna Shahab (University of Edinburgh), Mingcan Zhu (University of Edinburgh), Artemiy Margaritov (University of Edinburgh), and Boris Grot (University of Edinburgh)</i>	

5B: Measurement, Modeling, and Simulation

- Leveraging CPU Electromagnetic Emanations for Voltage Noise Characterization .573.....
Zacharias Hadjilambrou (University of Cyprus), Shidhartha Das (ARM Research, United Kingdom), Marco A. Antoniadis (University of Cyprus), and Yiannakis Sazeides (University of Cyprus)
- RpStacks-MT: A High-Throughput Design Evaluation Methodology for Multi-Core Processors .586....
Hanhwi Jang (POSTECH, South Korea), Jae-Eon Jo (POSTECH, South Korea), Jaewon Lee (Seoul National University, South Korea), and Jangwoo Kim (Seoul National University, South Korea)
- The EH Model: Early Design Space Exploration of Intermittent Processor Architectures .600.....
Joshua San Miguel (University of Wisconsin - Madison), Karthik Ganesan (University of Toronto), Mario Badr (University of Toronto), Chunqiu Xia (University of Toronto), Rose Li (University of Toronto), Hsuan Hsiao (University of Toronto), and Natalie Enright Jerger (University of Toronto)
- CounterMiner: Mining Big Performance Data from Hardware Counters .613.....
Yirong Lv (Shenzhen Institutes of Advanced Technology, Chinese Academy of Science, Shenzhen, China), Bin Sun (Capital Normal University, Beijing, China), Qingyi Luo (University of Southern California), Jing Wang (Capital Normal University), Zhibin Yu (Shenzhen Institute of Advanced Technology, Chinese Academy of Science, Shenzhen, China), and Xuehai Qian (University of Southern California, LA, USA)
- Taming the Killer Microsecond .627.....
Shenghsun Cho (Stony Brook University), Amoghavarsha Suresh (Stony Brook University), Tapti Palit (Stony Brook University), Michael Ferdman (Stony Brook University), and Nima Honarmand (Stony Brook University)

6A: Near Memory Computing

- Adaptive Scheduling for Systems with Asymmetric Memory Hierarchies .641.....
Po-An Tsai (Massachusetts Institute of Technology), Changping Chen (Massachusetts Institute of Technology), and Daniel Sanchez (Massachusetts Institute of Technology)
- Processing-in-Memory for Energy-Efficient Neural Network Training: A Heterogeneous Approach .655.....
Jiawen Liu (University of California, Merced), Hengyu Zhao (University of California, San Diego), Matheus A. Ogleari (University of California, Santa Cruz), Dong Li (University of California, Merced), and Jishen Zhao (University of California, San Diego)
- LerGAN: A Zero-Free, Low Data Movement and PIM-Based GAN Architecture .669.....
Haiyu Mao (Tsinghua University, China), Mingcong Song (University of Florida, United States), Tao Li (University of Florida, United States), Yuting Dai (Guizhou University, China), and Jiwu Shu (Tsinghua University, China)

Multi-dimensional Parallel Training of Winograd Layer on Memory-Centric Architecture .682.....	<i>Byungchul Hong (KAIST, Korea), Yeonju Ro (KAIST, Korea), and John Kim (KAIST, Korea)</i>
SCOPE: A Stochastic Computing Engine for DRAM-Based In-Situ Accelerator .696.....	<i>Shuangchen Li (University of California, Sanata Barbara), Alvin Oliver Glova (University of California, Sanata Barbara), Xing Hu (University of California, Sanata Barbara), Peng Gu (University of California, Sanata Barbara), Dimin Niu (Samsung), Krishna T. Malladi (Samsung), Hongzhong Zheng (Samsung), Bob Brennan (Samsung), and Yuan Xie (University of California, Sanata Barbara)</i>

6B: Near Memory Computing

Exploring and Optimizing Chipkill-Correct for Persistent Memory Based on High-Density NVRAMs .710.....	<i>Da Zhang (Virginia Tech), Vilas Sridharan (Advanced Micro Devices (AMD)), and Xun Jian (Virginia Tech)</i>
Comprehensive Evaluation of Supply Voltage Underscaling in FPGA on-Chip Memories .724.....	<i>Behzad Salami (Barcelona Supercomputing Center (BSC)), Osman S. Unsal (Barcelona Supercomputing Center (BSC)), and Adrian Cristal Kestelman (Barcelona Supercomputing Center (BSC))</i>
Error Correlation Prediction in Lockstep Processors for Safety-Critical Systems .737.....	<i>Emre Ozer (Arm Ltd.), Balaji Venu (Arm Ltd.), Xabier Iturbe (Arm Ltd.), Shidhartha Das (Arm Ltd.), Spyros Lyberis (Arm Ltd.), John Biggs (Arm Ltd.), Peter Harrod (Arm Ltd.), and John Penton (Arm Ltd.)</i>
Fault Site Pruning for Practical Reliability Analysis of GPGPU Applications .749.....	<i>Bin Nie (College of William and Mary, USA), Lishan Yang (College of William and Mary, USA), Adwait Jog (College of William and Mary, USA), and Evgenia Smirni (College of William and Mary, USA)</i>
SwapCodes: Error Codes for Hardware-Software Cooperative GPU Pipeline Error Detection .762.....	<i>Michael B. Sullivan (NVIDIA), Siva Kumar Sastry Hari (NVIDIA), Brian Zimmer (NVIDIA), Timothy Tsai (NVIDIA), and Stephen W. Keckler (NVIDIA)</i>

7: Best Papers

CEASER: Mitigating Conflict-Based Cache Attacks via Encrypted-Address and Remapping .775.....	<i>Moinuddin K. Qureshi (Georgia Institute of Technology)</i>
PipeProof: Automated Memory Consistency Proofs for Microarchitectural Specifications .788.....	<i>Yatin A. Manerkar (Princeton University, USA), Daniel Lustig (NVIDIA), Margaret Martonosi (Princeton University, USA), and Aarti Gupta (Princeton University, USA)</i>

- Application-Transparent Near-Memory Processing Architecture with Memory Channel Network .802.
Mohammad Alian (University of Illinois Urbana-Champaign), Seung Won Min (University of Illinois Urbana-Champaign), Hadi Asgharimoghaddam (University of Illinois Urbana-Champaign), Ashutosh Dhar (University of Illinois Urbana-Champaign), Dong Kai Wang (University of Illinois Urbana-Champaign), Thomas Roewer (IBM Research), Adam McPadden (IBM Systems), Oliver O'Halloran (IBM Systems), Deming Chen (University of Illinois Urbana-Champaign), Jinjun Xiong (IBM Research), Daehoon Kim (DGIST), Wen-mei Hwu (University of Illinois Urbana-Champaign), and Nam Sung Kim (University of Illinois Urbana-Champaign)
- End-to-End Automated Exploit Generation for Validating the Security of Processor Designs .815.....
Rui Zhang (University of North Carolina at Chapel Hill), Calvin Deutschbein (University of North Carolina at Chapel Hill), Peng Huang (Johns Hopkins University), and Cynthia Sturton (University of North Carolina at Chapel Hill)

8A: Non-Conventional Architectures

- Magic-State Functional Units: Mapping and Scheduling Multi-Level Distillation Circuits for Fault-Tolerant Quantum Architectures .828.....
Yongshan Ding (University of Chicago, USA), Adam Holmes (University of Chicago, USA), Ali Javadi-Abhari (IBM T.J. Watson Research Center, USA), Diana Franklin (University of Chicago, USA), Margaret Martonosi (Princeton University, USA), and Frederic Chong (University of Chicago, USA)
- MDACache: Caching for Multi-Dimensional-Access Memories .841.....
sumitha george (Pennsylvania State University), Minli Julie Liao (Pennsylvania State University), Huaipan Jiang (Pennsylvania State University), Jagadish B. Kotra (Pennsylvania State University, AMD Research), Mahmut T. Kandemir@cse.psu.edu, Jack Sampson (Pennsylvania State University), and Vijaykrishnan Narayanan (Pennsylvania State University)
- GeneSys: Enabling Continuous Learning through Neural Network Evolution in Hardware .855.....
Ananda Samajdar (Georgia Institute of Technology), Parth Mannan (Georgia Institute of Technology), Kartikay Garg (Georgia Institute of Technology), and Tushar Krishna (Georgia Institute of Technology)

8B: Mobile and Embedded Architectures

- CritICs Critiquing Criticality in Mobile Apps .867.....
Prasanna Venkatesh Rengasamy (The Pennsylvania State University), Haibo Zhang (The Pennsylvania State University), Shulin Zhao (The Pennsylvania State University), Nachiappan Chidambaram Nachiappan (IEEE Member), Anand Sivasubramaniam (The Pennsylvania State University), Mahmut T Kandemir (The Pennsylvania State University), and Chita R Das (The Pennsylvania State University)

EMPROF: Memory Profiling Via EM-Emanation in IoT and Hand-Held Devices .881.....	
	<i>Moumita Dey (Georgia Institute of Technology, USA), Alireza Nazari (Georgia Institute of Technology, USA), Alenka Zajic (Georgia Institute of Technology, USA), and Milos Prvulovic (Georgia Institute of Technology, USA)</i>
MAVBench: Micro Aerial Vehicle Benchmarking .894.....	
	<i>Behzad Boroujerdian (University of Texas, Austin), Hasan Genc (University of California, Berkeley), Srivatsan Krishnan (Harvard University), Wenzhi Cui (University of Texas, Austin), Aleksandra Faust (Google Brain), and Vijay Reddi (Harvard University)</i>

9A: Domain-Specific Architectures

Architectural Support for Efficient Large-Scale Automata Processing .908.....	
	<i>Hongyuan Liu (College of William & Mary), Mohamed Ibrahim (College of William & Mary), Onur Kayiran (Advanced Micro Devices, Inc.), Sreepathi Pai (University of Rochester), and Adwait Jog (College of William & Mary)</i>
ASPEN: A Scalable In-SRAM Architecture for Pushdown Automata .921.....	
	<i>Kevin Angstadt (University of Michigan, USA), Arun Subramaniyan (University of Michigan, USA), Elaheh Sadredini (University of Virginia, USA), Reza Rahimi (University of Virginia, USA), Kevin Skadron (University of Virginia, USA), Westley Weimer (University of Michigan, USA), and Reetuparna Das (University of Michigan)</i>
Morph: Flexible Acceleration for 3D CNN-Based Video Understanding .933.....	
	<i>Kartik Hegde (University of Illinois at Urbana-Champaign), Rohit Agrawal (University of Illinois at Urbana-Champaign), Yulun Yao (University of Illinois at Urbana-Champaign), and Christopher W Fletcher (University of Illinois at Urbana-Champaign)</i>

9B: Security-II

CheckMate: Automated Synthesis of Hardware Exploits and Security Litmus Tests .947.....	
	<i>Caroline Trippel (Princeton University), Daniel Lustig (NVIDIA), and Margaret Martonosi (Princeton University)</i>
Shadow Block: Accelerating ORAM Accesses with Data Duplication .961.....	
	<i>Xian Zhang (Peking University), Guangyu Sun (Peking University), Peichen Xie (Peking University), Chao Zhang (Peking University), Yannan Liu (The Chinese University of Hong Kong), Lingxiao Wei (The Chinese University of Hong Kong), Qiang Xu (The Chinese University of Hong Kong), and Chun Jason Xue (City University of Hong Kong)</i>
DAWG: A Defense Against Cache Timing Attacks in Speculative Execution Processors .974.....	
	<i>Vladimir Kiriansky (Massachusetts Institute of Technology, USA), Ilia Lebedev (Massachusetts Institute of Technology, USA), Saman Amarasinghe (Massachusetts Institute of Technology, USA), Srinivas Devadas (Massachusetts Institute of Technology, USA), and Joel Emer (NVIDIA / Massachusetts Institute of Technology, USA)</i>

Author Index ㉨㉨