

# **2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2020)**

**Athens, Greece  
17-21 October 2020**

**Pages 1-580**



**IEEE Catalog Number: CFP20071-POD  
ISBN: 978-1-7281-7384-9**

**Copyright © 2020 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:	CFP20071-POD
ISBN (Print-On-Demand):	978-1-7281-7384-9
ISBN (Online):	978-1-7281-7383-2

**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

# 2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO) **MICRO 2020**

## Table of Contents

<b>Message from the MICRO 2020 General Chair</b> .xvi.....	
<i>Dimitris Gizopoulos (National and Kapodistrian University of Athens)</i>	
<b>Message from the MICRO 2020 Program Co-Chairs</b> .xviii.....	
<i>Jun Yang (University of Pittsburgh) and Mattan Erez (University of Texas at Austin)</i>	
<b>Organizing Committee</b> .xxiii.....	
<b>Program Committee</b> .xxv.....	
<b>External Review Committee</b> .xxviii.....	
<b>Steering Committee</b> .xxxii.....	
<b>Additional External Reviewers</b> .xxxiii.....	
<b>MICRO-53 Sponsors</b> .xxxiv.....	
<b>Keynotes</b> .xxxv.....	

## Session 1A: Security and Privacy I

<b>Graphene: Strong yet Lightweight Row Hammer Protection</b> .1.....	
<i>Yeonhong Park (Seoul National University), Woosuk Kwon (Seoul National University), Eojin Lee (Seoul National University), Tae Jun Ham (Seoul National University), Jung Ho Ahn (Seoul National University), and Jae W. Lee (Seoul National University)</i>	
<b>Persist Level Parallelism: Streamlining Integrity Tree Updates for Secure Persistent Memory</b> .14.....	
<i>Alexander Freij (North Carolina State University), Shougang Yuan (North Carolina State University), Huiyang Zhou (North Carolina State University), and Yan Solihin (University of Central Florida)</i>	
<b>PThammer: Cross-User-Kernel-Boundary Rowhammer through Implicit Accesses</b> .28.....	
<i>Zhi Zhang (University of New South Wales, Australia and Data61, CSIRO, Australia), Yueqiang Cheng (Baidu Security), Dongxi Liu (Data61, CSIRO, Australia), Surya Nepal (Data61, CSIRO, Australia), Zhi Wang (Florida State University, America), and Yuval Yarom (University of Adelaide and Data61, CSIRO, Australia)</i>	

Draco: Architectural and Operating System Support for System Call Security .42.....  
*Dimitrios Skarlatos (University of Illinois at Urbana-Champaign),  
 Qingrong Chen (University of Illinois at Urbana-Champaign), Jianyan  
 Chen (University of Illinois at Urbana-Champaign), Tianyin Xu  
 (University of Illinois at Urbana-Champaign), and Josep Torrellas  
 (University of Illinois at Urbana-Champaign)*

## Session 1B: Machine Learning Accelerators with New Technologies

SuperNPU: An Extremely Fast Neural Processing Unit Using Superconducting Logic Devices .58....  
*Koki Ishida (Kyushu University), Ilkwon Byun (Seoul National  
 University), Ikki Nagaoka (Nagoya University), Kosuke Fukumitsu  
 (Kyushu University), Masamitsu Tanaka (Nagoya University), Satoshi  
 Kawakami (Kyushu University), Teruo Tanimoto (Kyushu University),  
 Takatsugu Ono (Kyushu University), Jangwoo Kim (Seoul National  
 University), and Koji Inoue (Kyushu University)*

Printed Machine Learning Classifiers .73.....  
*Muhammad Husnain Mubarik (University of Illinois Urbana-Champaign),  
 Dennis D. Weller (Karlsruhe Institute of Technology), Nathaniel Bleier  
 (University of Illinois Urbana-Champaign), Matthew Tomei (University  
 of Illinois Urbana-Champaign), Jasmin Aghassi-Hagmann (University of  
 Applied Sciences Offenburg), Mehdi B. Tahoori (Karlsruhe Institute of  
 Technology), and Rakesh Kumar (University of Illinois  
 Urbana-Champaign)*

Look-Up Table based Energy Efficient Processing in Cache Support for Neural Network  
 Acceleration .88.....  
*Akshay Krishna Ramanathan (Pennsylvania State University), Gurpreet S  
 Kalsi (Processor Architecture Research Lab, Intel Labs), Srivatsa  
 Srinivasa (Pennsylvania State University), Tarun Makesh Chandran  
 (Pennsylvania State University), Kamlesh R Pillai (Processor  
 Architecture Research Lab, Intel Labs), Om J Omer (Processor  
 Architecture Research Lab, Intel Labs), Vijaykrishnan Narayanan  
 (Pennsylvania State University), and Sreenivas Subramoney (Processor  
 Architecture Research Lab, Intel Labs)*

FReaC Cache: Folded-Logic Reconfigurable Computing in the Last Level Cache .102.....  
*Ashutosh Dhar (University of Illinois, Urbana Champaign), Xiaohao Wang  
 (University of Illinois, Urbana Champaign), Hubertus Franke (IBM  
 Research), Jinjun Xiong (IBM Research), Jian Huang (University of  
 Illinois, Urbana Champaign), Wen-mei Hwu (University of Illinois,  
 Urbana Champaign), Nam Sung Kim (University of Illinois, Urbana  
 Champaign), and Deming Chen (University of Illinois, Urbana Champaign)*

## Session 1C: Microarchitecture I

BranchNet: A Convolutional Neural Network to Predict Hard-to-Predict Branches .118.....  
*Siavash Zangeneh (University of Texas at Austin), Stephen Pruett  
 (University of Texas at Austin), Sangkug Lym (Nvidia), and Yale N.  
 Patt (University of Texas at Austin)*

CHiRP: Control-Flow History Reuse Prediction .131.....	
	<i>Samira Mirbagher-Ajorpaz (Texas A&amp;M University), Elba Garza (Texas A&amp;M University), Gilles Pokam (Intel Labs), and Daniel A. Jimenez (Texas A&amp;M University)</i>
I-SPY: Context-Driven Conditional Instruction Prefetching with Coalescing .146.....	
	<i>Tanvir Ahmed Khan (University of Michigan), Akshitha Sriraman (University of Michigan), Joseph Devietti (University of Pennsylvania), Gilles Pokam (Intel Corporation), Heiner Litz (University of California, Santa Cruz), and Baris Kasikci (University of Michigan)</i>
Improving the Utilization of Micro-Operation Caches in x86 Processors .160.....	
	<i>Jagdish B. Kotra (AMD Research) and John Kalamatianos (AMD Research)</i>

## Session 2A: Quantum Computing

Virtualized Logical Qubits: A 2.5D Architecture for Error-Corrected Quantum Computing .173.....	
	<i>Casey Duckering (University of Chicago), Jonathan M. Baker (University of Chicago), David I. Schuster (University of Chicago), and Frederic T. Chong (University of Chicago)</i>
Optimized Quantum Compilation for Near-Term Algorithms with OpenPulse .186.....	
	<i>Pranav Gokhale (University of Chicago), Ali Javadi-Abhari (IBM), Nathan Earnest (IBM), Yunong Shi (University of Chicago), and Frederic T. Chong (University of Chicago)</i>
Systematic Crosstalk Mitigation for Superconducting Qubits via Frequency-Aware Compilation.201	
	<i>Yongshan Ding (University of Chicago), Pranav Gokhale (University of Chicago), Sophia Fuhui Lin (University of Chicago), Richard Rines (University of Chicago), Thomas Propson (University of Chicago), and Frederic T. Chong (University of Chicago)</i>
Circuit Compilation Methodologies for Quantum Approximate Optimization Algorithm .215.....	
	<i>Mahabubul Alam (Pennsylvania State University), Abdullah Ash-Saki (Pennsylvania State University), and Swaroop Ghosh (Pennsylvania State University)</i>

## Session 2B: Robust Machine Learning

Fast-BCNN: Massive Neuron Skipping in Bayesian Convolutional Neural Networks .229.....	
	<i>Qiyu Wan (University of Houston) and Xin Fu (University of Houston)</i>
Ptolemy: Architecture Support for Robust Deep Learning .241.....	
	<i>Yiming Gan (University of Rochester), Yuxian Qiu (Shanghai Jiao Tong University), Jingwen Leng (Shanghai Jiao Tong University), Minyi Guo (Shanghai Jiao Tong University), and Yuhao Zhu (University of Rochester)</i>
Non-Blocking Simultaneous Multithreading: Embracing the Resiliency of Deep Neural Networks.256	
	<i>Gil Shomron (Technion - Israel Institute of Technology) and Uri Weiser (Technion - Israel Institute of Technology)</i>

Fidelity: Efficient Resilience Analysis Framework for Deep Learning Accelerators .270.....  
*Yi He (University of Chicago), Prasanna Balaprakash (Argonne National Laboratory), and Yanjing Li (University of Chicago)*

## Session 2C: Memory I

Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics .282.....  
*Minesh Patel (ETH Zürich), Jeremie S. Kim (Carnegie Mellon University, ETH Zürich), Taha Shahroodi (ETH Zürich), Hasan Hassan (ETH Zürich), and Onur Mutlu (ETH Zürich, Carnegie Mellon University)*

DStress: Automatic Synthesis of DRAM Reliability Stress Viruses using Genetic Algorithms .298...  
*Lev Mukhanov (Queen's University Belfast), Dimitrios S. Nikolopoulos (Virginia Tech, USA), and Georgios Karakonstantis (Queen's University Belfast)*

FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching ...  
313

*Yaohua Wang (National University of Defense Technology), Lois Orosa (ETH Zürich), Xiangjun Peng (Chinese University of Hong Kong / National University of Defense Technology), Yang Guo (National University of Defense Technology), Saugata Ghose (University of Illinois at Urbana-Champaign / Carnegie Mellon University), Minesh Patel (ETH Zürich), Jeremie S. Kim (ETH Zürich), Juan Gómez Luna (ETH Zürich), Mohammad Sadrosadati (Institute of Research in Fundamental Sciences), Nika Mansouri Ghiasi (ETH Zürich), and Onur Mutlu (ETH Zürich / Carnegie Mellon University)*

PerpLE: Improving the Speed and Effectiveness of Memory Consistency Testing .329.....  
*Themis Melissaris (Princeton University), Markos Markakis (Princeton University), Kelly Shaw (Williams College), and Margaret Martonosi (Princeton University)*

## Session 3A: Near/In-Memory Computing

CATCAM: Constant-Time Alteration Ternary CAM with Scalable In-Memory Architecture .342....  
*Dibei Chen (Tsinghua University), Zhaoshi Li (Tsinghua University), Tianzhu Xiong (Southeast University), Zhiwei Liu (Tsinghua University), Jun Yang (Southeast University), Shouyi Yin (Tsinghua University), Shaojun Wei (Tsinghua University), and Leibo Liu (Tsinghua University)*

DUAL: Acceleration of Clustering Algorithms using Digital-Based Processing In-Memory .356.....  
*Mohsen Imani (University of California Irvine), Saikishan Pampana (University of California San Diego), Saransh Gupta (University of California San Diego), Minxuan Zhou (University of California San Diego), Yeseong Kim (Daegu Institute of Science and Technology), and Tajana Rosing (University of California San Diego)*

Newton: A DRAM-Maker's Accelerator-in-Memory (AiM) Architecture for Machine Learning .372  
*Mingxuan He (Purdue University), Choungki Song (SK Hynix Inc.), Ilkon Kim (SK Hynix Inc.), Chunseok Jeong (SK Hynix Inc.), Seho Kim (Sk Hynix Inc.), Il Park (SK Hynix Inc.), Mithuna Thottethodi (Purdue University), and T. N. Vijaykumar (Purdue University)*

AQUOMAN: An Analytic-Query Offloading Machine .386.....  
*Shuotao Xu (MIT CSAIL), Thomas Bourgeat (MIT CSAIL), Tianhao Huang (MIT CSAIL), Hojun Kim (DGIST), Sungjin Lee (DGIST), and Arvind Arvind (MIT CSAIL)*

MOUSE: Inference In Non-Volatile Memory for Energy Harvesting Applications .400.....  
*Salonik Resch (University of Minnesota), S. Karen Khatamifard (University of Minnesota), Zamshed I. Chowdhury (University of Minnesota), Masoud Zabihi (University of Minnesota), Zhengyang Zhao (University of Minnesota), Husrev Cilasun (University of Minnesota), Jian-Ping Wang (University of Minnesota), Sachin S. Sapatnekar (University of Minnesota), and Ulya R. Karpuzcu (University of Minnesota)*

## Session 3B: Compilation, Modeling, and Simulation

More with Less - Deriving More Translation Rules with Less Training Data for DBTs Using Parameterization .415.....  
*Jinhu Jiang (Fudan University), Rongchao Dong (Fudan University), Zhongjun Zhou (Fudan University), Changheng Song (Fudan University), Wenwen Wang (Department of Computer Science, University of Georgia), Pen-Chung Yew (Department of Computer Science and Engineering, University of Minnesota, Twin Cities), and Weihua Zhang (Fudan University)*

Optimizing the Memory Hierarchy by Compositing Automatic Transformations on Computations and Data .427.....  
*Jie Zhao (State Key Laboratory of Mathematical Engineering and Advanced Computing) and Peng Di (Huawei Technologies Co. , Ltd.)*

DiffTune: Optimizing CPU Simulator Parameters with Learned Differentiable Surrogates .442.....  
*Alex Renda (MIT CSAIL), Yishen Chen (MIT CSAIL), Charith Mendis (MIT CSAIL), and Michael Carbin (MIT CSAIL)*

Predicting Execution Times With Partial Simulations in Virtual Memory Research: Why and How .456.....  
*Mohammad Agbarya (Technion - Israel Institute of Technology), Idan Yaniv (Technion - Israel Institute of Technology), Jayneel Gandhi (VMware Research), and Dan Tsafir (Technion - Israel Institute of Technology & VMware Research)*

gem5-SALAM: A System Architecture for LLVM-Based Accelerator Modeling .471.....  
*Samuel Rogers (UNC Charlotte), Joshua Slycord (UNC Charlotte), Mohammedreza Baharani (UNC Charlotte), and Hamed Tabkhi (UNC Charlotte)*

## Session 3C: Non-volatile Memories

- Shaving Retries with Sentinels for Fast Read over High-Density 3D Flash .483.....  
*Qiao Li (Department of Computer Science, City University of Hong Kong), Min Ye (Department of Computer Science, City University of Hong Kong), Yufei Cui (Department of Computer Science, City University of Hong Kong), Liang Shi (East China Normal University), Xiaoqiang Li (YEESTOR Microelectronics Co., Ltd), Tei-Wei Kuo (Department of Computer Science, City University of Hong Kong), and Chun Jason Xue (Department of Computer Science, City University of Hong Kong)*
- Characterizing and Modeling Non-Volatile Memory Systems .496.....  
*Zixuan Wang (University of California, San Diego), Xiao Liu (University of California, San Diego), Jian Yang (University of California, San Diego), Theodore Michailidis (University of California, San Diego), Steven Swanson (University of California, San Diego), and Jishen Zhao (University of California, San Diego)*
- P-INSPECT: Architectural Support for Programmable Non-Volatile Memory Frameworks .509.....  
*Apostolos Kokolis (University of Illinois at Urbana-Champaign), Thomas Shull (University of Illinois at Urbana-Champaign), Jian Huang (University of Illinois at Urbana-Champaign), and Josep Torrellas (University of Illinois at Urbana-Champaign)*
- Unbounded Hardware Transactional Memory for a Hybrid DRAM/NVM Memory System .525.....  
*Jungi Jeong (Purdue University), Jaewan Hong (KAIST), Seungryoul Maeng (KAIST), Changhee Jung (Purdue University), and Youngjin Kwon (KAIST)*
- (Almost) Fence-Less Persist Ordering .539.....  
*Sara Mahdizadeh Shahri (Pennsylvania State University), Seyed Armin Vakil Ghahani (Pennsylvania State University), and Aasheesh Kolli (Pennsylvania State University)*

## Session 4A: Microarchitecture II

- Speculative Enforcement of Store Atomicity .555.....  
*Alberto Ros (University of Murcia) and Stefanos Kaxiras (Uppsala University)*
- Boosting Store Buffer Efficiency with Store-Prefetch Bursts .568.....  
*Juan M. Cebrian (University of Murcia), Stefanos Kaxiras (Uppsala University), and Alberto Ros (University of Murcia)*
- D-SOAP: Dynamic Spatial Orientation Affinity Prediction for Caching in Multi-Orientation Memory Systems .581.....  
*Minli Julie Liao (Department of Computer Science and Engineering, School of Electrical Engineering and Computer Science, Pennsylvania State University, Pennsylvania, USA) and Jack Sampson (Department of Computer Science and Engineering, School of Electrical Engineering and Computer Science, Pennsylvania State University, Pennsylvania, USA)*
- Pipette: Improving Core Utilization on Irregular Applications through Intra-Core Pipeline Parallelism .596.....  
*Quan M. Nguyen (MIT CSAIL) and Daniel Sanchez (MIT CSAIL)*



RnR: A Software-Assisted Record-and-Replay Hardware Prefetcher .609.....  
Chao Zhang (Lehigh University), Yuan Zeng (Lehigh University), John  
Shalf (Lawrence Berkeley National Laboratory), and Xiaochen Guo  
(Lehigh University)

## Session 4B: Resource Management

ConfuciuX: Autonomous Hardware Resource Assignment for DNN Accelerators using  
Reinforcement Learning .622.....  
Sheng-Chun Kao (Georgia Institute of Technology), Geonhwa Jeong  
(Georgia Institute of Technology), and Tushar Krishna (Georgia  
Institute of Technology)

Gemini: Learning to Manage CPU Power for Latency-Critical Search Engines .637.....  
Liang Zhou (University of California, Riverside), Laxmi N. Bhuyan  
(University of California, Riverside), and K. K. Ramakrishnan  
(University of California, Riverside)

CuttleSys: Data-Driven Resource Management for Interactive Services on Reconfigurable  
Multicores .650.....  
Neeraj Kulkarni (Microsoft), Gonzalo Gonzalez-Pumariiega (Cornell  
University), Amulya Khurana (Cornell University), Christine A.  
Shoemaker (National University of Singapore), Christina Delimitrou  
(Cornell University), and David H. Albonesi (Cornell University)

Jumanji: The Case for Dynamic NUCA in the Datacenter .665.....  
Brian C. Schwedock (Carnegie Mellon University) and Nathan Beckmann  
(Carnegie Mellon University)

Planaria: Dynamic Architecture Fission for Spatial Multi-Tenant Acceleration of Deep  
Neural Networks .681.....  
Soroush Ghodrati (University of California, San Diego), Byung Hoon Ahn  
(University of California, San Diego), Joon Kyung Kim (University of  
California, San Diego), Sean Kinzer (University of California, San  
Diego), Brahmendra Reddy Yatham (University of California, San Diego),  
Navateja Alla (University of California, San Diego), Hardik Sharma  
(Bigstream Inc.), Mohammad Alian (Kansas University), Eiman Ebrahimi  
(NVIDIA Research), Nam Sung Kim (University of Illinois  
Urbana-Champaign), Cliff Young (Google Inc.), and Hadi Esmaeilzadeh  
(University of California, San Diego)

## Session 4C: Machine Learning Accelerators I

VR-DANN: Real-Time Video Recognition via Decoder-Assisted Neural Network Acceleration .698  
Zhuoran Song (Shanghai Jiao Tong University), Feiyang Wu (Shanghai  
Jiao Tong University), Xueyuan Liu (Shanghai Jiao Tong University),  
Jing Ke (Shanghai Jiao Tong University, Biren Research), Naifeng Jing  
(Shanghai Jiao Tong University, Biren Research), and Xiaoyao Liang  
(Shanghai Jiao Tong University, Biren Research)

- Procrustes: A Dataflow and Accelerator for Sparse Deep Neural Network Training .711.....  
*Dingqing Yang (The University of British Columbia), Amin Ghasemazar (The University of British Columbia), Xiaowei Ren (The University of British Columbia), Maximilian Golub (Microsoft Corporation), Guy Lemieux (The University of British Columbia), and Mieszko Lis (The University of British Columbia)*
- Duplo: Lifting Redundant Memory Accesses of Deep Neural Networks for GPU Tensor Cores .725  
*Hyeonjin Kim (Yonsei University), Sungwoo Ahn (Yonsei University), Yunho Oh (École Polytechnique Fédérale de Lausanne (EPFL)), Bogil Kim (Yonsei University), Won Woo Ro (Yonsei University), and William J. Song (Yonsei University)*
- DUET: Boosting Deep Neural Network Efficiency on Dual-Module Architecture .738.....  
*Liu Liu (University of California, Santa Barbara), Zheng Qu (University of California, Santa Barbara), Lei Deng (University of California, Santa Barbara), Fengbin Tu (University of California, Santa Barbara), Shuangchen Li (University of California, Santa Barbara), Xing Hu (University of California, Santa Barbara), Zhenyu Gu (Alibaba DAMO Academy), Yufei Ding (University of California, Santa Barbara), and Yuan Xie (University of California, Santa Barbara)*

## Session 5A: Machine Learning Accelerators II

- TFE: Energy-Efficient Transferred Filter-Based Engine to Compress and Accelerate Convolutional Neural Networks .751.....  
*Huiyu Mo (Tsinghua University), Leibo Liu (Tsinghua University), Wenjing Hu (Tsinghua University), Wenping Zhu (Tsinghua University), Qiang Li (Intel Corporation), Ang Li (Tsinghua University), Shouyi Yin (Tsinghua University), Jian Chen (Alibaba Group), Xiaowei Jiang (Alibaba Group), and Shaojun Wei (Tsinghua University)*
- MatRaptor: A Sparse-Sparse Matrix Multiplication Accelerator Based on Row-Wise Product .766...  
*Nitish Srivastava (Cornell University), Hanchen Jin (Cornell University), Jie Liu (Cornell University), David Albonesei (Cornell University), and Zhiru Zhang (Cornell University)*
- TensorDash: Exploiting Sparsity to Accelerate Deep Neural Network Training .781.....  
*Mostafa Mahmoud (University of Toronto), Isak Edo (University of Toronto), Ali Hadi Zadeh (University of Toronto), Omar Mohamed Awad (University of Toronto), Gennady Pekhimenko (University of Toronto, Vector Institute), Jorge Albericio (Cerebras Systems), and Andreas Moshovos (University of Toronto, Vector Institute)*
- SAVE: Sparsity-Aware Vector Engine for Accelerating DNN Training and Inference on CPUs .796..  
*Zhangxiaowen Gong (University of Illinois at Urbana-Champaign; Intel Labs), Houxiang Ji (University of Illinois at Urbana-Champaign), Christopher W. Fletcher (University of Illinois at Urbana-Champaign), Christopher J. Hughes (Intel Labs), Sara Baghsorkhi (Intel Labs), and Josep Torrellas (University of Illinois at Urbana-Champaign)*
- GOBO: Quantizing Attention-Based NLP Models for Low Latency and Energy Efficient Inference.811  
*Ali Hadi Zadeh (University of Toronto), Isak Edo (University of Toronto), Omar Mohamed Awad (University of Toronto), and Andreas Moshovos (University of Toronto)*

## Session 5B: Cloud and Datacenter

TrainBox: An Extreme-Scale Neural Network Training Server Architecture by Systematically Balancing Operations .825.....	
<i>Pyeongsu Park (Seoul National University), Heetaek Jeong (Seoul National University), and Jangwoo Kim (Seoul National University)</i>	
Coordinated Priority-Aware Charging of Distributed Batteries in Oversubscribed Data Centers .839.....	
<i>Sulav Malla (Facebook, Inc.; University of South Florida), Qingyuan Deng (Facebook, Inc.), Zoh Ebrahimzadeh (Facebook, Inc.), Joe Gasperetti (Facebook, Inc.), Sajal Jain (Facebook, Inc.), Parimala Kondety (Facebook, Inc.), Thiara Ortiz (Facebook, Inc.), and Debra Vieira (Facebook, Inc.)</i>	
HyperPlane: A Scalable Low-Latency Notification Accelerator for Software Data Planes .852.....	
<i>Amirhossein Mirhosseini (University of Michigan), Hossein Golestani (University of Michigan), and Thomas F. Wensch (University of Michigan)</i>	
ThymesisFlow: A Software-Defined, HW/SW Co-Designed Interconnect Stack for Rack-Scale Memory Disaggregation .868.....	
<i>Christian Pinto (IBM Research Europe), Dimitris Syriovelis (IBM Research Europe), Michele Gazzetti (IBM Research Europe), Panos Koutsovasilis (IBM Research Europe), Andrea Reale (IBM Research Europe), Kostas Katrinis (IBM Research Europe), and H. Peter Hofstee (IBM Systems)</i>	
A Benchmarking Framework for Interactive 3D Applications in the Cloud .881.....	
<i>Tianyi Liu (The University of Texas at San Antonio), Sen He (The University of Texas at San Antonio), Sunzhou Huang (The University of Texas at San Antonio), Danny Tsang (The University of Texas at San Antonio), Lingjia Tang (University of Michigan, Ann Arbor), Jason Mars (University of Michigan, Ann Arbor), and Wei Wang (The University of Texas at San Antonio)</i>	

## Session 5C: Domain-Specific Architecture

- A Locality-Aware Energy-Efficient Accelerator for Graph Mining Applications .895.....  
*Pengcheng Yao (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Long Zheng (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Zhen Zeng (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Yu Huang (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Chuangyi Gui (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Xiaofei Liao (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), Hai Jin (National Engineering Research Center for Big Data Technology and System/Services Computing Technology and System Lab/Cluster and Grid Computing Lab, Huazhong University of Science and Technology, China), and Jingling Xue (UNSW Sydney, Australia)*
- GraphPulse: An Event-Driven Hardware Accelerator for Asynchronous Graph Processing .908.....  
*Shafiur Rahman (University of California, Riverside), Nael Abu-Ghazaleh (University of California, Riverside), and Rajiv Gupta (University of California, Riverside)*
- AWB-GCN: A Graph Convolutional Network Accelerator with Runtime Workload Rebalancing .922  
*Tong Geng (Boston University), Ang Li (Pacific Northwest National Laboratory), Runbin Shi (The University of Hong Kong), Chunshu Wu (Boston University), Tianqi Wang (Boston University), Yanfei Li (Zhejiang University), Pouya Haghi (Boston University), Antonino Tumeo (Pacific Northwest National Laboratory), Shuai Che (Microsoft), Steve Reinhardt (Microsoft), and Martin C. Herbordt (Boston University)*
- SeedEx: A Genome Sequencing Accelerator for Optimal Alignments in Subminimal Space .937.....  
*Daichi Fujiki (University of Michigan), Shunhao Wu (University of Michigan), Nathan Ozog (University of Michigan), Kush Goliya (University of Michigan), David Blaauw (University of Michigan), Satish Narayanasamy (University of Michigan), and Reetuparna Das (University of Michigan)*

GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis .951.....

*Damla Senol Cali (Carnegie Mellon University), Gurpreet S. Kalsi (Processor Architecture Research Lab, Intel Labs), Zülal Bingöl (Bilkent University), Can Firtina (ETH Zürich), Lavanya Subramanian (Facebook), Jeremie S. Kim (ETH Zurich, Carnegie Mellon University), Rachata Ausavarungnirun (King Mongkut’s University of Technology North Bangkok), Mohammed Alser (ETH Zürich), Juan Gomez-Luna (ETH Zürich), Amirali Boroumand (Carnegie Mellon University), Anant Nori (Processor Architecture Research Lab, Intel Labs), Allison Scibisz (Carnegie Mellon University), Sreenivas Subramoney (Processor Architecture Research Lab, Intel Labs), Can Alkan (Bilkent University), Saugata Ghose (University of Illinois at Urbana–Champaign, Carnegie Mellon University), and Onur Mutlu (ETH Zürich, Carnegie Mellon University, Bilkent University)*

## Session 6A: GPGPU

Selective Replication in Memory-Side GPU Caches .967.....

*Xia Zhao (Academy of Military Science), Magnus Jahre (Norwegian University of Science and Technology, Norway), and Lieven Eeckhout (Ghent University)*

Deterministic Atomic Buffering .981.....

*Yuan Hsi Chou (University of British Columbia), Christopher Ng (University of British Columbia), Shaylin Cattell (University of British Columbia), Jeremy Intan (University of Wisconsin), Matthew D. Sinclair (University of Wisconsin, AMD Research), Joseph Devietti (University of Pennsylvania), Timothy G. Rogers (Purdue University), and Tor M. Aamodt (University of British Columbia)*

BOW: Breathing Operand Windows to Exploit Bypassing in GPUs .996.....

*Hodjat Asghari Esfeden (University of California Riverside), Amirali Abdolrashidi (University of California Riverside), Shafiur Rahman (University of California Riverside), Daniel Wong (University of California Riverside), and Nael Abu-Ghazaleh (University of California Riverside)*

MDM: The GPU Memory Divergence Model .1009.....

*Lu Wang (Ghent University, Belgium), Magnus Jahre (Norwegian University of Science and Technology), Almutaz Adileh (Huawei, Toga Networks), and Lieven Eeckhout (Ghent University, Belgium)*

Locality-Centric Data and Threadblock Management for Massive GPUs .1022.....

*Mahmoud Khairy (Purdue University), Vadim Nikiforov (Purdue University), David Nellans (NVIDIA), and Timothy G. Rogers (Purdue University)*

## Session 6B: Mobile and Embedded Architecture

Mesorasi: Architecture Support for Point Cloud Analytics via Delayed-Aggregation .....	1037
<i>Yu Feng (University of Rochester), Boyuan Tian (University of Rochester), Tiancheng Xu (University of Rochester), Paul Whatmough (Arm Research), and Yuhao Zhu (University of Rochester)</i>	
FlexWatts: A Power- and Workload-Aware Hybrid Power Delivery Network for Energy-Efficient Microprocessors .....	1051
<i>Jawad Haj-Yahya (ETH Zürich), Mohammed Alser (ETH Zürich), Jeremie S. Kim (ETH Zürich), Lois Orosa (ETH Zürich), Efraim Rotem (Intel), Avi Mendelson (Technion; Nanyang Technological University), Anupam Chattopadhyay (Nanyang Technological University), and Onur Mutlu (ETH Zürich)</i>	
Building the Computing System for Autonomous Micromobility Vehicles: Design Constraints and Architectural Optimizations .....	1067
<i>Bo Yu (PerceptIn), Wei Hu (PerceptIn), Leimeng Xu (PerceptIn), Jie Tang (South China University of Technology), Shaoshan Liu (PerceptIn), and Yuhao Zhu (University of Rochester)</i>	
AutoScale: Energy Efficiency Optimization for Stochastic Edge Inference Using Reinforcement Learning .....	1082
<i>Young Geun Kim (Arizona State University) and Carole-Jean Wu (Arizona State University, Facebook AI)</i>	
NCPU: An Embedded Neural CPU Architecture on Resource-Constrained Low Power Devices for Real-Time End-to-End Performance .....	1097
<i>Tianyu Jia (Northwestern Univeristy), Yuhao Ju (Northwestern Univeristy), Russ Joseph (Northwestern Univeristy), and Jie Gu (Northwestern Univeristy)</i>	

## Session 6C: Security and Privacy II

CaSA: End-to-end Quantitative Security Analysis of Randomly Mapped Caches .....	1110
<i>Thomas Bourgeat (MIT), Jules Drean (MIT), Yuheng Yang (MIT/University of Chinese Academy of Science), Lillian Tsai (MIT), Joel Emer (Nvidia/MIT), and Mengjia Yan (MIT)</i>	
PerSpectron: Detecting Invariant Footprints of Microarchitectural Attacks with Perceptron .....	1124
<i>Samira Mirbagher-Ajorpaz (Texas A&amp;M University), Gilles Pokam (Intel Labs), Esmail Mohammadian-Koruyeh (University of California, Riverside), Elba Garza (Texas A&amp;M University), Nael Abu-Ghazaleh (University of California, Riverside), and Daniel A. Jimenez (Texas A&amp;M University)</i>	
Speculation Invariance (InvarSpec): Faster Safe Execution through Program Analysis .....	1138
<i>Zirui Neil Zhao (University of Illinois at Urbana-Champaign), Houxiang Ji (University of Illinois at Urbana-Champaign), Mengjia Yan (Massachusetts Institute of Technology), Jiyong Yu (University of Illinois at Urbana-Champaign), Christopher W. Fletcher (University of Illinois at Urbana-Champaign), Adam Morrison (Tel Aviv University), Darko Marinov (University of Illinois at Urbana-Champaign), and Josep Torrellas (University of Illinois at Urbana-Champaign)</i>	
Hardware-Based Always-On Heap Memory Safety .....	1153
<i>Yonghae Kim (Georgia Institute of Technology), Jaekyu Lee (Arm Research), and Hyesoon Kim (Georgia Institute of Technology)</i>	

## Author Index