

2021 IEEE 39th International Conference on Computer Design (ICCD 2021)

**Virtual Conference
24 – 27 October 2021**



**IEEE Catalog Number: CFP21ICD-POD
ISBN: 978-1-6654-3220-7**

**Copyright © 2021 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:	CFP21ICD-POD
ISBN (Print-On-Demand):	978-1-6654-3220-7
ISBN (Online):	978-1-6654-3219-1
ISSN:	1063-6404

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

2021 IEEE 39th International Conference on Computer Design (ICCD) **ICCD 2021**

Table of Contents

Message from the General Chairs	xvii
Organizing Committee	xviii
Message from the Program Chairs	xix
Program Committee	xx
Special Session PC	xxiv
Keynotes	xxv
Tutorials	xxvii

Special Session: Cross-Layer Approximations for System-Level Optimizations

Special Session: ADAPT: ANN-Controlled System-Level Runtime Adaptable APproximate CompuTing	1
<i>Prattay Chowdhury (The University of Texas at Dallas, USA) and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
Special Session: How Much Quality is Enough Quality? A Case for Acceptability in Approximate Designs	5
<i>Isaiás Felzmann (University of Campinas, Brazil), João Fabrício Filho (University of Campinas, Brazil; Federal University of Technology, Brazil), Juliane Regina de Oliveira (University of Campinas, Brazil), and Lucas Wanner (University of Campinas, Brazil)</i>	
Special Session: When Dataflows Converge: Reconfigurable and Approximate Computing for Emerging Neural Networks	9
<i>Di Wu (University of Wisconsin–Madison, USA) and Joshua San Miguel (University of Wisconsin–Madison, USA)</i>	
Special Session: Approximate TinyML Systems: Full System Approximations for Extreme Energy-Efficiency in Intelligent Edge Devices	13
<i>Arnab Raha (Edge.AI Advanced Architecture Group, VPU IP, Intel Corporation, USA), Soumendu Ghosh (Purdue University, USA), Debabrata Mohapatra (Edge.AI Advanced Architecture Group, VPU IP, Intel Corporation, USA), Deepak A. Mathaikutty (Edge.AI Advanced Architecture Group, VPU IP, Intel Corporation, USA), Raymond Sung (Edge.AI Advanced Architecture Group, VPU IP, Intel Corporation, USA), Cormac Brick (Edge.AI Advanced Architecture Group, VPU IP, Intel Corporation, USA), and Vijay Raghunathan (Purdue University, USA)</i>	

Best Papers Session

Flipping Bits to Share Crossbars in ReRAM-Based DNN Accelerator	17
<i>Lei Zhao (University of Pittsburgh, USA), Youtao Zhang (University of Pittsburgh, USA), and Jun Yang (University of Pittsburgh, USA)</i>	
Block-LSM: An Ether-Aware Block-Ordered LSM-Tree Based Key-Value Storage Engine	25
<i>Zehao Chen (Shandong University, China), Bingzhe Li (Oklahoma State University, USA), Xiaojun Cai (Shandong University, China), Zhiping Jia (Shandong University, China), Zhaoyan Shen (Shandong University, China), Yi Wang (Shenzhen University, China), and Zili Shao (The Chinese University of Hong Kong, Hong Kong)</i>	
DynPaC: Coarse-Grained, Dynamic, and Partially Reconfigurable Array for Streaming Applications	33
<i>Cheng Tan (Pacific Northwest National Laboratory, USA), Tong Geng (Pacific Northwest National Laboratory, USA), Chenhao Xie (Pacific Northwest National Laboratory, USA), Nicolas Bohm Agostini (Pacific Northwest National Laboratory, USA), Jiajia Li (Pacific Northwest National Laboratory, USA), Ang Li (Pacific Northwest National Laboratory, USA), Kevin Barker (Pacific Northwest National Laboratory, USA), and Antonino Tumeo (Pacific Northwest National Laboratory, USA)</i>	
NIST-Lite: Randomness Testing of RNGs on an Energy-Constrained Platform	41
<i>Cheng-Yen Lee (Texas A&M University, USA), Kunal Bharathi (Texas A&M University, USA), Joellen Lansford (Texas A&M University, USA), and Sunil P. Khatri (Texas A&M University, USA)</i>	
Seer-SSD: Bridging Semantic Gap Between Log-Structured File Systems and SSDs to Reduce SSD Write Amplification	49
<i>You Zhou (Huazhong University of Science and Technology, China), Ke Wang (Huazhong University of Science and Technology, China), Fei Wu (Huazhong University of Science and Technology, China), Changsheng Xie (Huazhong University of Science and Technology, China), and Hao Lv (Huazhong University of Science and Technology, China)</i>	
Novel Ultra-Low-Voltage Flip-Flops : Near-V _{th} Modeling and VLSI Integration	57
<i>Anuradha Ranasinghe (University of Twente, Netherlands) and Sabih H. Gerez (University of Twente, Netherlands)</i>	

Session 1A: Novel Architectures

Flexible Instruction Set Architecture for Programmable Look-up Table Based Processing-in-Memory	66
<i>Mark Connolly (Rochester Institute of Technology, USA), Purab Ranjan Sutradhar (Rochester Institute of Technology, USA), Mark Indovina (Rochester Institute of Technology, USA), and Amlan Ganguly (Rochester Institute of Technology, USA)</i>	

Stochastic Iterative Approximation: Software/hardware Techniques for Adjusting Aggressiveness of Approximation	74
<i>Tomoki Nakamura (The University of Tokyo, Japan), Kazutaka Tomida (The University of Tokyo, Japan), Shouta Kouno (The University of Tokyo, Japan), Hidetsugu Irie (The University of Tokyo, Japan), and Shuichi Sakai (The University of Tokyo, Japan)</i>	
The Accuracy and Efficiency of Posit Arithmetic	83
<i>Stefan Dan Ciocirlan (National University of Singapore, Singapore; University Politehnica of Bucharest, Romania), Dumitrel Loghin (National University of Singapore, Singapore), Lavanya Ramapantulu (Independent Researcher), Nicolae Tapus (University Politehnica of Bucharest, Romania), and Yong Meng Teo (National University of Singapore, Singapore)</i>	
Accurate and Fast Performance Modeling of Processors with Decoupled Front-end	88
<i>Yuya Degawa (The University of Tokyo, Japan), Toru Koizumi (The University of Tokyo, Japan), Tomoki Nakamura (The University of Tokyo, Japan), Ryota Shioya (The University of Tokyo, Japan), Junichiro Kadomoto (The University of Tokyo, Japan), Hidetsugu Irie (The University of Tokyo, Japan), and Shuichi Sakai (The University of Tokyo, Japan)</i>	
AdaptBit-HD: Adaptive Model Bitwidth for Hyperdimensional Computing	93
<i>Justin Morris (University of California San Diego, USA; San Diego State University, USA), Si Thu Kaung Set (University of California San Diego, USA), Gadi Rosen (University of California San Diego, USA), Mohsen Imani (University of California Irvine, USA), Baris Aksanli (San Diego State University, USA), and Tajana Rosing (University of California San Diego, USA)</i>	
MasterMind: Many-Accelerator SoC Architecture for Real-Time Brain-Computer Interfaces	101
<i>Guy Eichler (Columbia University, USA), Luca Piccolboni (Columbia University, USA), Davide Giri (Columbia University, USA), and Luca P. Carloni (Columbia University, USA)</i>	

Session 1B: Arithmetic Designs and Analyses

An Efficient Hybrid Parallel Compression Approximate Multiplier	109
<i>Shangshang Yao (National University of Defense Technology, China), Liang Zhang (National University of Defense Technology, China), Qiong Wang (National University of Defense Technology, China), and Li Shen (National University of Defense Technology, China)</i>	
Run-Time Configurable Approximate Multiplier using Significance-Driven Logic Compression	117
<i>Ibrahim Haddadi (Newcastle University, UK; TAIBAH University, Kingdom of Saudi Arabia), Issa Qiqieh (Al-Balqa Applied University, Jordan), Rishad Shafik (Newcastle University, UK), Fei Xia (Newcastle University, UK), Mohammed Al-hayanni (The University of Technology, Iraq), and Alex Yakovlev (Newcastle University, UK)</i>	
A Comprehensive Exploration of the Parallel Prefix Adder Tree Space	125
<i>Teodor-Dumitru Ene (Oklahoma State University, USA) and James E. Stine (Oklahoma State University, USA)</i>	

CROP: FPGA Implementation of High-Performance Polynomial Multiplication in Saber KEM Based on Novel Cyclic-Row Oriented Processing Strategy	130
<i>Jiafeng Xie (Villanova University, USA), Pengzhou He (Villanova University, USA), and Chiou-Yng Lee (Lunghwa University of Science & Technology, Taiwan)</i>	
Energy-Efficient MAC Units for Fused Posit Arithmetic	138
<i>Raul Murillo (Complutense University of Madrid, Spain), David Mallasén (Complutense University of Madrid, Spain), Alberto A. Del Barrio (Complutense University of Madrid, Spain), and Guillermo Botella (Complutense University of Madrid, Spain)</i>	

Session 2A: Persistent Memory

NRHI: A Concurrent Non-Rehashing Hash Index for Persistent Memory	146
<i>Xinyu Li (SKLCA, ICT, CAS, China), Huimin Cui (SKLCA, ICT, CAS, China), and Lei Liu (SKLCA, ICT, CAS, China)</i>	
HASDH: A Hotspot-Aware and Scalable Dynamic Hashing for Hybrid DRAM-NVM Memory	154
<i>Zhengtao Li (HuaZhong University of Science and Technology, China), Zhipeng Tan (HuaZhong University of Science and Technology, China), and Jianxi Chen (HuaZhong University of Science and Technology, China)</i>	
EFM: Elastic Flash Management to Enhance Performance of Hybrid Flash Memory	162
<i>Bingzhe Li (Oklahoma State University, USA), Bo Yuan (Rutgers University, USA), and David Du (University of Minnesota, USA)</i>	
Dynamic File Cache Optimization for Hybrid SSDs with High-Density and Low-Cost Flash Memory	170
<i>Ben Gu (East China Normal University, China), Longfei Luo (East China Normal University, China), Yina Lv (East China Normal University, China), Changlong Li (East China Normal University, China), and Liang Shi (Ministry of Education, China; East China Normal University, China)</i>	
Consistent RDMA-Friendly Hashing on Remote Persistent Memory	174
<i>Xinxin Liu (Huazhong University of Science and Technology, China), Yu Hua (Huazhong University of Science and Technology, China), and Rong Bai (Huazhong University of Science and Technology, China)</i>	

Session 2B: Test and Verification

Low-Cost Sequential Logic Circuit Design Considering Single Event Double-Node Upsets and Single Event Transients	178
<i>Ramin Rajaei (University of Notre Dame, USA), Michael Niemier (University of Notre Dame, USA), and Xiaobo Sharon Hu (University of Notre Dame, USA)</i>	
Rectification of Integer Arithmetic Circuits using Computer Algebra Techniques	186
<i>Vikas Rao (University of Utah, USA), Haden Ondricek (University of Utah, USA), Priyank Kalla (University of Utah, USA), and Florian Enescu (Georgia State University, USA)</i>	

Differential Testing of x86 Instruction Decoders with Instruction Operand Inferring Algorithm	196
<i>Guang Wang (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences), Ziyuan Zhu (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences), Shuan Li (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences), Xu Cheng (Peking University), and Dan Meng (Chinese Academy of Sciences, China; University of Chinese Academy of Sciences)</i>	
Comprehensive Failure Analysis Against Soft Errors from Hardware and Software Perspectives....	204
<i>Yohan Ko (Yonsei University, Republic of Korea), Hwisoo So (Yonsei University, Republic of Korea), Jinhyo Jung (Yonsei University, Republic of Korea), Kyoungwoo Lee (Yonsei University, Republic of Korea), and Aviral Shrivastava (Arizona State University, USA)</i>	
MetaTableLite: An Efficient Metadata Management Scheme for Tagged-Pointer-Based Memory Safety	208
<i>Dongwei Chen (Peking University), Dong Tong (Peking University), Chun Yang (Peking University), and Xu Cheng (Peking University)</i>	
HammerFilter: Robust Protection and Low Hardware Overhead Method for RowHammer	212
<i>Kwangrae Kim (Hanyang University, South Korea), Jeonghyun Woo (University of Illinois at Urbana-Champaign, USA), Junsu Kim (Hanyang University, South Korea), and Ki-Seok Chung (Hanyang University, South Korea)</i>	

Session 3A: Non Volatile Memory

SimiEncode: A Similarity-Based Encoding Scheme to Improve Performance and Lifetime of Non-Volatile Main Memory	220
<i>Suzhen Wu (Xiamen University, China), Jiapeng Wu (Xiamen University, China), Zhirong Shen (Xiamen University, China), Zhihao Zhang (Xiamen University, China), Zuocheng Wang (Xiamen University, China), and Bo Mao (Xiamen University, China)</i>	
Accelerating Sub-Block Erase in 3D NAND Flash Memory	228
<i>Hongbin Gong (Xiamen University), Zhirong Shen (Xiamen University), and Jiwu Shu (Xiamen University; Tsinghua University)</i>	
Understanding and Optimizing Hybrid SSD with High-Density and Low-Cost Flash Memory	236
<i>Liang Shi (East China Normal University, China; Huazhong University of Science and Technology, China), Longfei Luo (East China Normal University, China), Yina Lv (East China Normal University, China), Shicheng Li (East China Normal University, China), Changlong Li (East China Normal University, China), and Edwin Hsing-Mean Sha (East China Normal University, China)</i>	
WAS-Deletion: Workload-Aware Secure Deletion Scheme for Solid-State Drives	244
<i>Bingzhe Li (Oklahoma State University, USA) and David Du (University of Minnesota, USA)</i>	
Write-Optimized and Consistent RDMA-Based Non-Volatile Main Memory Systems	248
<i>Xinxin Liu (Huazhong University of Science and Technology, China), Yu Hua (Huazhong University of Science and Technology, China), Xuan Li (Huazhong University of Science and Technology, China), and Qifan Liu (Huazhong University of Science and Technology, China)</i>	

Session 3B: Hardware Accelerators

An Ultra-Efficient Look-up Table Based Programmable Processing in Memory Architecture for Data Encryption	252
<i>Purab Ranjan Sutradhar (Rochester Institute of Technology, USA), Kanad Basu (University of Texas at Dallas, USA), Sai Manoj Pudukotai Dinakarrao (George Mason University, USA), and Amlan Ganguly (Rochester Institute of Technology, USA)</i>	
ReSpar: Reordering Algorithm for ReRAM-Based Sparse Matrix-Vector Multiplication Accelerator	260
<i>Yi-Jou Hsiao (Academia Sinica, Taiwan), Chin-Fu Nien (Academia Sinica, Taiwan), and Hsiang-Yun Cheng (Academia Sinica, Taiwan)</i>	
Fast and Low-Cost Mitigation of ReRAM Variability for Deep Learning Applications	269
<i>Sugil Lee (Ulsan National Institute of Science and Technology (UNIST), South Korea), Mohammed Fouda (University of California-Irvine, USA), Jongeun Lee (Ulsan National Institute of Science and Technology (UNIST), South Korea), Ahmed Eltawil (University of California-Irvine, USA; King Abdulla University of Science and Technology, Saudi Arabia), and Fadi Kurdahi (University of California-Irvine, USA)</i>	
WiNN: Wireless Interconnect Based Neural Network Accelerator	277
<i>Siqin Liu (Ohio University), Sushanth Karmunchi (California State University Fresno), Avinash Karanth (Ohio University), Soumyasanta Laha (California State University Fresno), and Savas Kaya (Ohio University)</i>	
Universal Neural Network Acceleration via Real-Time Loop Blocking	285
<i>Jiaqi Zhang (University of Florida), Xiangru Chen (University of Florida), and Sandip Ray (University of Florida)</i>	

Session 4A: Exploiting Parallelism

Exploiting Intra-SM Parallelism in GPUs via Persistent and Elastic Blocks	290
<i>Han Zhao (Shanghai Jiao Tong University, China), Weihao Cui (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Jieru Zhao (Shanghai Jiao Tong University, China), Jingwen Leng (Shanghai Jiao Tong University, China), and Minyi Guo (Shanghai Jiao Tong University, China)</i>	
Improved GPU Implementations of the Pair-HMM Forward Algorithm for DNA Sequence Alignment	299
<i>Enliang Li (University of Illinois at Urbana-Champaign, USA), Subho S. Banerjee (University of Illinois at Urbana-Champaign, USA), Sitao Huang (University of Illinois at Urbana-Champaign, USA), Ravishankar K. Iyer (University of Illinois at Urbana-Champaign, USA), and Deming Chen (University of Illinois at Urbana-Champaign, USA)</i>	

CHARM: Collaborative Host and Accelerator Resource Management for GPU Datacenters	307
<i>Wei Zhang (Shanghai Jiao Tong University Shanghai, China), Kaihua Fu (Shanghai Jiao Tong University Shanghai, China), Ningxin Zheng (Microsoft research Asia Shanghai, China), Quan Chen (Shanghai Jiao Tong University Shanghai, China), Chao Li (Shanghai Jiao Tong University Shanghai, China), Wenli Zheng (Shanghai Jiao Tong University Shanghai, China), and Minyi Guo (Shanghai Jiao Tong University Shanghai, China)</i>	
Empirical Guide to Use of Persistent Memory for Large-Scale In-Memory Graph Analysis	316
<i>Hanyeoreum Bae (Korea Advanced Institute of Science and Technology), Miryeong Kwon (Korea Advanced Institute of Science and Technology), Donghyun Gouk (Korea Advanced Institute of Science and Technology), Sanghyun Han (Korea Advanced Institute of Science and Technology), Sungjoon Koh (Korea Advanced Institute of Science and Technology), Changrim Lee (Korea Advanced Institute of Science and Technology), Dongchul Park (Sookmyung Women's University), and Myoungsoo Jung (Korea Advanced Institute of Science and Technology)</i>	
Stochastic-HD: Leveraging Stochastic Computing on Hyper-Dimensional Computing	321
<i>Yilun Hao (University of California San Diego, USA), Sransh Gupta (University of California San Diego, USA), Justin Morris (University of California San Diego, USA; San Diego State University, USA), Behnam Khaleghi (University of California San Diego, USA), Baris Aksanli (San Diego State University, USA), and Tajana Rosing (University of California San Diego, USA)</i>	
HyperData: A Data Transfer Accelerator for Software Data Planes Based on Targeted Prefetching	326
<i>Hossein Golestani (University of Michigan) and Thomas F. Wenisch (University of Michigan)</i>	

Session 4B: Potpourri

UH-JLS: A Parallel Ultra-High Throughput JPEG-LS Encoding Architecture for Lossless Image Compression	335
<i>Xuan Wang (University of Science and Technology of China, China), Lei Gong (University of Science and Technology of China, China), Chao Wang (University of Science and Technology of China, China), Xi Li (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)</i>	
PRL: Standardizing Performance Monitoring Library for High-Integrity Real-Time Systems	344
<i>Jeremy Giesen (Barcelona Supercomputing Center; Universitat Politècnica de Catalunya), Enrico Mezzetti (Barcelona Supercomputing Center), Jaume Abella (Barcelona Supercomputing Center), and Francisco J. Cazorla (Barcelona Supercomputing Center)</i>	
CIDAN: Computing in DRAM with Artificial Neurons	349
<i>Gian Singh (Arizona State University, USA), Ankit Wagle (Arizona State University, USA), Sarma Vrudhula (Arizona State University, USA), and Sunil Khatri (Texas A&M University, USA)</i>	

Resonance-Based Power-Efficient Pulse Generator Design with Corresponding Distribution Network	357
<i>Ke Jia (State Key Laboratory of Computer Architecture, ICT, CAS; University of Chinese Academy of Sciences), Liang Yang (Loongson Technology Co., Ltd.), Jian Wang (State Key Laboratory of Computer Architecture, ICT, CAS; University of Chinese Academy of Sciences), Bin Lin (State Key Laboratory of Computer Architecture, ICT, CAS), Hao Wang (Loongson Technology Co., Ltd.), and RuiKai Shi (State Key Laboratory of Computer Architecture, ICT, CAS)</i>	
Reconfigurable Array for Analog Applications	361
<i>Ziyi Chen (Drexel University) and Ioannis Savidis (Drexel University)</i>	
CoRe-ECO: Concurrent Refinement of Detailed Place-and-Route for an Efficient ECO Automation	366
<i>Chung-Kuan Cheng (University of California San Diego, USA), Andrew B. Kahng (University of California San Diego, USA), Ilgweon Kang (Digital & Signoff Group, Cadence Design Systems Inc., USA), Minsoo Kim (University of California San Diego, USA), Daeyeal Lee (University of California San Diego, USA), Bill Lin (University of California San Diego, USA), Dongwon Park (University of California San Diego, USA), and Mingyu Woo (University of California San Diego, USA)</i>	

Session 5A: High Performance Systems

Efficient Table-Based Polynomial on FPGA	374
<i>Marco Barbone (Imperial College London, UK), Bas W. Kwaadgras (Maxeler Technologies, UK), Uwe Oelfke (Joint Department of Physics at The Institute of Cancer Research and The Royal Marsden NHS Foundation Trust, UK), Wayne Luk (Imperial College London, UK), and Georgi Gaydadjiev (University of Groningen, Netherlands; Imperial College London, UK)</i>	
Legion: Tailoring Grouped Neural Execution Considering Heterogeneity on Multiple Edge Devices	383
<i>Kyunghwan Choi (Hanyang University, Korea), Seongju Lee (Hanyang University, Korea), Beom Woo Kang (Hanyang University, Korea), and Yongjun Park (Hanyang University, Korea)</i>	
Premier: A Concurrency-Aware Pseudo-Partitioning Framework for Shared Last-Level Cache	391
<i>Xiaoyang Lu (Illinois Institute of Technology), Rujia Wang (Illinois Institute of Technology), and Xian-He Sun (Illinois Institute of Technology)</i>	
Chopin: Composing Cost-Effective Custom Chips with Algorithmic Chiplets	395
<i>Pete Ehrett (University of Michigan), Todd Austin (University of Michigan), and Valeria Bertacco (University of Michigan)</i>	
Compiling and Optimizing Real-World Programs for STRAIGHT ISA	400
<i>Toru Koizumi (The University of Tokyo, Japan), Shu Sugita (The University of Tokyo, Japan), Ryota Shioya (The University of Tokyo, Japan), Junichiro Kadomoto (The University of Tokyo, Japan), Hidetsugu Irie (The University of Tokyo, Japan), and Shuichi Sakai (The University of Tokyo, Japan)</i>	

GVSoc: A Highly Configurable, Fast and Accurate Full-Platform Simulator for RISC-V Based IoT Processors	409
<i>Nazareno Bruschi (University of Bologna, Italy), Germain Haugou (ETH, Switzerland; GreenWaves-Technologies, France), Giuseppe Tagliavini (University of Bologna, Italy), Francesco Conti (University of Bologna, Italy), Luca Benini (University of Bologna, Italy; ETH, Switzerland), and Davide Rossi (University of Bologna, Italy)</i>	

Session 5B: Memory Design and Security

SME: ReRAM-Based Sparse-Multiplication-Engine to Squeeze-Out Bit Sparsity of Neural Network	417
<i>Fangxin Liu (Shanghai Jiao Tong University, China; Shanghai Qi Zhi Institute, China), Wenbo Zhao (Shanghai Jiao Tong University, China; Shanghai Qi Zhi Institute, China), Zhezhi He (Shanghai Jiao Tong University, China), Zongwu Wang (Shanghai Jiao Tong University, China), Yilong Zhao (Shanghai Jiao Tong University, China), Tao Yang (Shanghai Jiao Tong University, China), Jingnai Feng (Shanghai Jiao Tong University, China), Xiaoyao Liang (Shanghai Jiao Tong University, China), and Li Jiang (Shanghai Jiao Tong University, China; Shanghai Qi Zhi Institute, China; Shanghai Jiao Tong University, China)</i>	
Improving the Heavy Re-Encryption Overhead of Split Counter Mode Encryption for NVM	425
<i>Qi Pei (The State University of New York at Binghamton, USA) and Seunghee Shin (The State University of New York at Binghamton, USA)</i>	
Discreet-PARA: Rowhammer Defense with Low Cost and High Efficiency	433
<i>Yicheng Wang (University of Illinois at Chicago, US), Yang Liu (University of Illinois at Chicago, US), Peiyun Wu (University of Illinois at Chicago, US), and Zhao Zhang (University of Illinois at Chicago, US)</i>	
Conciliating Speed and Efficiency on Cache Compressors	442
<i>Daniel Rodrigues Carvalho (Univ Rennes, Inria, CNRS, IRISA, France) and André Seznec (Univ Rennes, Inria, CNRS, IRISA, France)</i>	
POMI: Polling-Based Memory Interface for Hybrid Memory System	447
<i>Trung Le (University of Illinois at Chicago, USA), Zhao Zhang (University of Illinois at Chicago, USA), and Zhichun Zhu (University of Illinois at Chicago, USA)</i>	

Session 6A: Storage

Fault-Aware Prediction-Guided Page Offlining for Uncorrectable Memory Error Prevention	456
<i>Xiaoming Du (Intel Corporation, China), Cong Li (Intel Corporation, China), Shen Zhou (Intel Corporation, China), Xian Liu (ByteDance, China), Xiaohan Xu (ByteDance, China), Tianjiao Wang (ByteDance, China), and Shijian Ge (ByteDance, China)</i>	
A High-Performance Post-Deduplication Delta Compression Scheme for Packed Datasets	464
<i>Yucheng Zhang (Nanchang University, China), Hong Jiang (University of Texas at Arlington, USA), Mengtian Shi (Hubei University of Technology, China), Chunzhi Wang (Hubei University of Technology, China), Nan Jiang (East China Jiaotong University, China), and Xinyun Wu (Hubei University of Technology, China)</i>	

Erasure-Coded Multi-Block Updates Based on Hybrid Writes and Common XORs First	472
<i>Yujun Liu (Guangdong University of Technology, China), Bing Wei (Guangdong University of Technology, China), Jigang Wu (Guangdong University of Technology, China), and Limin Xiao (Beihang University, China)</i>	
PSACS: Highly-Parallel Shuffle Accelerator on Computational Storage	480
<i>Chen Zou (University of Chicago, USA), Hui Zhang (Samsung Semiconductor, Inc., USA), Andrew A. Chien (University of Chicago, USA), and Yang Seok Ki (Samsung Semiconductor, Inc., USA)</i>	
Intelligent Prediction of Flash Lifetime via Online Domain Adaptation	488
<i>Ruixiang Ma (Huazhong University of Science & Technology, China), Fei Wu (Huazhong University of Science & Technology, China), and Changsheng Xie (Huazhong University of Science & Technology, China)</i>	

Session 6B Design Automation and Analysis

Model Synthesis for Communication Traces of System Designs	492
<i>Hao Zheng (University of South Florida), Md Rubel Ahmed (University of South Florida), Parijat Mukherjee (Intel), Mahesh C. Ketkar (Intel), and Jin Yang (Intel)</i>	
T-TSP: Transient-Temperature Based Safe Power Budgeting in Multi-/Many-Core Processors	500
<i>Sobhan Niknam (University of Amsterdam, The Netherlands), Anuj Pathania (University of Amsterdam, The Netherlands), and Andy D. Pimentel (University of Amsterdam, The Netherlands)</i>	
HosNa: A DPC++ Benchmark Suite for Heterogeneous Architectures	509
<i>Najmeh Nazari Bavarsad (University of California Davis), Hosein Mohammadi Makrani (University of California Davis), Hossein Sayadi (California State University Long Beach), Lawrence Landis (Intel), Setareh Rafatirad (University of California Davis), and Houman Hodayoun (University of California Davis)</i>	
Optimal Transistor Placement Combined with Global In-Cell Routing in Standard Cell Layout Synthesis	517
<i>Kyeongrok Jo (Seoul National University, Korea) and Taewhan Kim (Seoul National University, Korea)</i>	
RDP^3: Rapid Domain Platform Performance Prediction for Design Space Exploration	525
<i>Jinghan Zhang (Northeastern University, USA), Mehrshad Zandigohar (Northeastern University, USA), and Gunar Schirner (Northeastern University, USA)</i>	

Session 7A Neural Networks

Smart-DNN: Efficiently Reducing the Memory Requirements of Running Deep Neural Networks on Resource-Constrained Platforms	533
<i>Zhenbo Hu (Harbin Institute of Technology, Shenzhen), Xiangyu Zou (Harbin Institute of Technology, Shenzhen), Wen Xia (Harbin Institute of Technology, Shenzhen; Peng Cheng Laboratory; Chinese Academy of Sciences), Yuhong Zhao (Chinese Academy of Sciences), Weizhe Zhang (Harbin Institute of Technology, Shenzhen; Cyberspace Security Research Center, Peng Cheng Laboratory), and Donglei Wu (Harbin Institute of Technology, Shenzhen)</i>	
QD-Compressor: a Quantization-Based Delta Compression Framework for Deep Neural Networks	542
<i>Shuyu Zhang (Harbin Institute of Technology, Shenzhen), Donglei Wu (Harbin Institute of Technology, Shenzhen), Haoyu Jin (Harbin Institute of Technology, Shenzhen), Xiangyu Zou (Harbin Institute of Technology, Shenzhen), Wen Xia (Harbin Institute of Technology, Shenzhen; Peng Cheng Laboratory; Chinese Academy of Sciences), and Xiaojia Huang (Harbin Institute of Technology, Shenzhen)</i>	
AGQFL: Communication-Efficient Federated Learning via Automatic Gradient Quantization in Edge Heterogeneous Systems	551
<i>Zirui Lian (University of Science and Technology of China, China), Jing Cao (University of Science and Technology of China, China), Yanru Zuo (University of Science and Technology of China, China), Weihong Liu (University of Science and Technology of China, China), and Zongwei Zhu (University of Science and Technology of China, China)</i>	
ModelShield: A Generic and Portable Framework Extension for Defending Bit-Flip Based Adversarial Weight Attacks	559
<i>Yanan Guo (University of Pittsburgh, USA), Liang Liu (University of Pittsburgh, USA), Yueqiang Cheng (NIO, USA), Youtao Zhang (University of Pittsburgh, USA), and Jun Yang (University of Pittsburgh, USA)</i>	
WidePipe: High-Throughput Deep Learning Inference System on a Cluster of Neural Processing Units	563
<i>Lixian Ma (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, China), En Shao (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China), Yueyuan Zhou (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, China), and Guangming Tan (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS, China; University of Chinese Academy of Sciences, China)</i>	
Exploiting Online Locality and Reduction Parallelism for Sampled Dense Matrix Multiplication on GPUs	567
<i>Zhongming Yu (Tsinghua University, China), Guohao Dai (Tsinghua University, China), Guyue Huang (UC Santa Barbara, USA), Yu Wang (Tsinghua University, China), and Huazhong Yang (Tsinghua University, China)</i>	

Session 7B Security and Trust

APT: Efficient Side-Channel Analysis Framework Against Inner Product Masking Scheme	575
<i>Jingdian Ming (Chinese Academy of Science, China; University of Chinese Academy of Science, China), Wei Cheng (Telecom Paris, Polytechnique de Paris, France), Yongbin Zhou (Chinese Academy of Science, China; University of Chinese Academy of Science, China; Nanjing University of Science and Technology, China), and Huizhong Li (Chinese Academy of Science, China; University of Chinese Academy of Science, China)</i>	
An Efficient Non-Profiled Side-Channel Attack on the CRYSTALS-Dilithium Post-Quantum Signature	583
<i>Zhaohui Chen (University of Chinese Academy of Sciences, China), Emre Karabulut (North Carolina State University, USA), Aydin Aysu (North Carolina State University, USA), Yuan Ma (CAS, China), and Jiwu Jing (University of Chinese Academy of Sciences, China)</i>	
Functional Locking Through Omission: From HLS to Obfuscated Design	591
<i>Zi Wang (The University of Texas at Dallas, USA), Shayan Omais Mohammed (The University of Texas at Dallas, USA), Yiorgos Makris (The University of Texas at Dallas, USA), and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
Security Analysis of State-of-the-art Scan Obfuscation Technique	599
<i>Yogendra Sao (Indian Institute of Technology Bhilai, India) and Sk Subidh Ali (Indian Institute of Technology Bhilai, India)</i>	
QFlow: Quantitative Information Flow for Security-Aware Hardware Design in Verilog	603
<i>Lennart M. Reimann (RWTH Aachen University, Germany), Luca Hanel (RWTH Aachen University, Germany), Dominik Sisejkovic (RWTH Aachen University, Germany), Farhad Merchant (RWTH Aachen University, Germany), and Rainer Leupers (RWTH Aachen University, Germany)</i>	
Efficient Methods for SoC Trust Validation using Information Flow Verification	608
<i>Khitam M. Alatoun (University of Cincinnati, USA), Shanmukha Murali Achyutha (University of Cincinnati, USA), and Ranga Vemuri (University of Cincinnati, USA)</i>	
Author Index	617