

2019 IEEE 37th International Conference on Computer Design (ICCD 2019)

**Abu Dhabi, United Arab Emirates
17 – 20 November 2019**



**IEEE Catalog Number: CFP19ICD-POD
ISBN: 978-1-7281-1215-2**

**Copyright © 2019 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:	CFP19ICD-POD
ISBN (Print-On-Demand):	978-1-7281-1215-2
ISBN (Online):	978-1-5386-6648-7
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

2019 IEEE 37th International Conference on Computer Design (ICCD) ICCD 2019

Table of Contents

Message from General Chairs and Program Chair	xviii
Organizing Committee	xx
Program Committee	xxi
Sub Reviewers	xxv
Sponsors	xxvii

Session 1A: Optimized Design Methodology

Reading Between the Dies: Cross-SLR Covert Channels on Multi-Tenant Cloud FPGAs	1
<i>Ilias Giechaskiel (University of Oxford, United Kingdom), Kasper Rasmussen (University of Oxford, United Kingdom), and Jakub Szefer (Yale University, USA)</i>	
IPSA: Integer Programming via Sparse Approximation for Efficient Test-Chip Design	11
<i>Qicheng Huang (Carnegie Mellon University), Chenlei Fang (Carnegie Mellon University), Zeye Liu (Carnegie Mellon University), Ruizhou Ding (Carnegie Mellon University), and R. D. Shawn Blanton (Carnegie Mellon University)</i>	
Stealthy Rootkits in Smart Grid Controllers	20
<i>Prashanth Krishnamurthy (NYU Tandon School of Engineering), Hossein Salehghaffari (NYU Tandon School of Engineering), Shiva Duraisamy (NYU Tandon School of Engineering), Ramesh Karri (NYU Tandon School of Engineering), and Farshad Khorrami (NYU Tandon School of Engineering)</i>	
Efficient Scalable Three Operand Multiplier Over GF(2 ^m) Based on Novel Decomposition Strategy	29
<i>Chiou-Yng Lee (Lunghwa University of Science & Technology) and Jiafeng Xie (Villanova University)</i>	

Session 1B: Accelerators and Machine Learning I

An FPGA Implementation of Stochastic Computing-Based LSTM	38
<i>Guy Maor (The University of Texas at Dallas, USA, U.S.), Xiaoming Zeng (The University of Texas at Dallas, USA, U.S.), Zhendong Wang (The University of Texas at Dallas, USA, U.S.), and Yang Hu (The University of Texas at Dallas, USA, U.S.)</i>	

Process Variation Mitigation on Convolutional Neural Network Accelerator Architecture	47
<i>Maodi Ma (Jilin University, China), Jingweijia Tan (Jilin University, China), Xiaohui Wei (Jilin University, China), and Kaige Yan (Jilin University, China)</i>	
Characterizing On-Chip Traffic Patterns in General-Purpose GPUs: A Deep Learning Approach	56
<i>Yunfan Li (Oregon State University), Drew Penney (Oregon State University), Abhishek Ramamurthy (Oregon State University), and Lizhong Chen (Oregon State University)</i>	
AccUDNN: A GPU Memory Efficient Accelerator for Training Ultra-Deep Neural Networks	65
<i>Jinrong Guo (Institute of Information Engineering, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Wantao Liu (Institute of Information Engineering, Chinese Academy of Sciences, China), Wang Wang (Institute of Information Engineering, Chinese Academy of Sciences, China; University of Chinese Academy of Sciences, China), Chunrong Yao (Nanjing University of Science and Technology, China), Jizhong Han (Institute of Information Engineering, Chinese Academy of Sciences, China), Ruixuan Li (Institute of Information Engineering, Chinese Academy of Sciences, China), Yijun Lu (Alibaba Cloud Computing Co. Ltd., China), and Songlin Hu (Institute of Information Engineering, Chinese Academy of Sciences, China)</i>	

Session 2A: Architectural Advances for IoT Applications

A Distributed Scheme for Accelerating Semantic Video Segmentation on An Embedded Cluster	73
<i>Hsuan-Kung Yang (National Tsing Hua University), Tsu-Jui Fu (National Tsing Hua University), Po-Han Chiang (National Tsing Hua University), Kuan-Wei Ho (National Tsing Hua University), and Chun-Yi Lee (National Tsing Hua University)</i>	
TiLA: Twin-in-the-Loop Architecture for Cyber-Physical Production Systems	82
<i>Heejong Park (Nanyang Technological University), Arvind Easwaran (Nanyang Technological University), and Sidharta Andalarn (Delta Electronics)</i>	
RRAMedy: Protecting ReRAM-Based Neural Network from Permanent and Soft Faults During Its Lifetime	91
<i>Wen Li (Institute of Computing Technology, Chinese Academy of Sciences, China), Ying Wang (Institute of Computing Technology, Chinese Academy of Sciences, China), Huawei Li (Institute of Computing Technology, Chinese Academy of Sciences, China), and Xiaowei Li (Institute of Computing Technology, Chinese Academy of Sciences, China)</i>	
WiXI: An Inter-Chip Wireless Bus Interface for Shape-Changeable Chiplet-Based Computers	100
<i>Junichiro Kadamoto (The University of Tokyo, Japan), Hidetsugu Irie (The University of Tokyo), and Shuichi Sakai (The University of Tokyo)</i>	

Session 2B: Architecture & Compilers II

Bandwidth-Aware Last-Level Caching: Efficiently Coordinating Off-Chip Read and Write Bandwidth	109
<i>Mainak Chaudhuri (Indian Institute of Technology Kanpur, India), Jayesh Gaur (Intel Processor Architecture Research Lab, India), and Sreenivas Subramoney (Intel Processor Architecture Research Lab, India)</i>	
To Update or Not To Update?: Bandwidth-Efficient Intelligent Replacement Policies for DRAM Caches	119
<i>Vinson Young (Georgia Institute of Technology) and Moinuddin K. Qureshi (Georgia Institute of Technology)</i>	
Compiler-Assisted Selection of Hardware Acceleration Candidates from Application Source Code	129
<i>Georgios Zacharopoulos (Università della Svizzera italiana, Switzerland), Lorenzo Ferretti (Università della Svizzera italiana, Switzerland), Giovanni Ansaloni (Università della Svizzera italiana, Switzerland), Giuseppe Di Guglielmo (Columbia University, USA), Luca Carloni (Columbia University, USA), and Laura Pozzi (Università della Svizzera italiana, Switzerland)</i>	
A New Traffic Offloading Method with Slow Switching Optical Device in Exascale Computer	138
<i>En Shao (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS), Guangming Tan (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS), Zhan Wang (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS), Guojun Yuan (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS), and Ninghui Sun (State Key Laboratory of Computer Architecture, Institute of Computing Technology, CAS)</i>	

Session 3A: Advances in the Design and Implementation of Neural Networks

ReRAM Crossbar-Based Analog Computing Architecture for Naive Bayesian Engine	147
<i>Bing Wu (Huazhong University of Science and Technology, China), Dan Feng (Huazhong University of Science and Technology, China), Wei Tong (Huazhong University of Science and Technology, China), Jingning Liu (Huazhong University of Science and Technology, China), Chengning Wang (Huazhong University of Science and Technology, China), Wei Zhao (Huazhong University of Science and Technology, China), and Mengye Peng (Huazhong University of Science and Technology, China)</i>	
HBUNN - Hybrid Binary-Unary Neural Network: Realizing a Complete CNN on an FPGA	156
<i>Sayed Abdolrasoul Faraji (University of Minnesota Twin Cities, Minneapolis, MN, US), Gaurav Singh (University of Minnesota Twin Cities, Minneapolis, MN, US), and Kia Bazargan (University of Minnesota Twin Cities, Minneapolis, MN, US)</i>	
Hardware Acceleration of Multilayer Perceptron Based on Inter-Layer Optimization	164
<i>Shenggang Chen (National University of Defense Technology, China) and Zhonghai Lu (KTH Royal Institute of Technology, Sweden)</i>	

NR-MPA: Non-Recovery Compression Based Multi-Path Packet-Connected-Circuit Architecture of Convolution Neural Networks Accelerator	173
<i>Gaoming Du (HeFei University of Technology, China), Zhenwen Yang (HeFei University of Technology, China), Zhenmin Li (HeFei University of Technology, China), Duoli Zhang (HeFei University of Technology, China), Yongsheng Yin (HeFei University of Technology, China), and Zhonghai Lu (KTH Royal Institute of Technology, Sweden)</i>	
A Memory-Access-Efficient Adaptive Implementation of kNN on FPGA through HLS	177
<i>Xiaojia Song (San Diego State University), Tao Xie (San Diego State University), and Stephen Fischer (Samsung Semiconductor, Inc.)</i>	

Session 3B: Storage Systems

Adaptive Write Interference Management with Efficient Mapping for Shingled Recording Disks	181
<i>Ming-Chang Lee (National Chiao Tung University, Taiwan), Li-Pin Chang (National Chiao Tung University, Taiwan), Sung-Ming Wu (National Chiao Tung University, Taiwan), and Wei-Shang Yui (National Chiao Tung University)</i>	
RepEC-Duet: Ensure High Reliability and Performance for Deduplicated and Delta-Compressed Storage Systems	190
<i>Chunxue Zuo (Huazhong University of Science and Technology, China), Fang Wang (Huazhong University of Science and Technology, China), Ping Huang (Temple University, USA), Yuchong Hu (Huazhong University of Science and Technology, China), and Dan Feng (Huazhong University of Science and Technology, China)</i>	
Lifelong Disk Failure Prediction via GAN-Based Anomaly Detection	199
<i>Tianming Jiang (Huazhong University of Science and Technology, China), Jiangfeng Zeng (Central China Normal University, China), Ke Zhou (Huazhong University of Science and Technology, China), Ping Huang (Temple University, USA), and Tianming Yang (Huanghuai University)</i>	
An Efficient and Flexible Metadata Management Layer for Local File Systems	208
<i>Yubo Liu (Sun Yat-sen University, China), Hongbo Li (Sun Yat-sen University, China), Yutong Lu (Sun Yat-sen University, China), Zhiguang Chen (Sun Yat-sen University, China), and Ming Zhao (Arizona State University, USA)</i>	

Session 4A: Innovation on Safety and Security for Robust Real-Time Systems

Security-Driven Codesign with Weakly-Hard Constraints for Real-Time Embedded Systems	217
<i>Hengyi Liang (Northwestern University, USA), Zhilu Wang (Northwestern University, USA), Debayan Roy (Technical University of Munich, Germany), Soumyajit Dey (Indian Institute of Technology, India), Samarjit Chakraborty (Technical University of Munich, Germany), and Qi Zhu (Northwestern University, USA)</i>	

Software Timing Analysis for Complex Hardware with Survivability and Risk Analysis	227
<i>Sergi Vilar dell (Universitat Politecnica de Catalunya, Barcelona Supercomputing Center), Isabel Serra (Centre de Recerca Matematica, Barcelona Supercomputing Center), Jaume Abella (Barcelona Supercomputing Center), Joan Del Castillo (Universitat Autonoma de Barcelona), and Francisco J. Cazorla (Barcelona Supercomputing Center)</i>	
Integrating Cyber-Attack Defense Techniques into Real-Time Cyber-Physical Systems	237
<i>Xiaochen Hao (Northeastern University, China), Mingsong Lv (Northeastern University, China), Jiesheng Zheng (Northeastern University, China), Zhengkui Zhang (Northeastern University, China), and Wang Yi (Uppsala University, Sweden)</i>	
Astraea: Self-Balancing Federated Learning for Improving Classification Accuracy of Mobile Deep Learning Applications	246
<i>Moming Duan (Chongqing University, China), Duo Liu (Chongqing University, China), Xianzhang Chen (Chongqing University, China), Yujuan Tan (Chongqing University, China), Jinting Ren (Chongqing University, China), Lei Qiao (Beijing Institute of Control Engineering, China), and Liang Liang (Chongqing University, China)</i>	

Session 4B: Processor and Memory Architectures

A Comparative Analysis on the Impact of Bank Contention in STT-MRAM and SRAM Based LLCs	255
<i>Timon Evenblij (imec, Belgium), Manu Perumkunnil (imec, Belgium), Francky Catthoor (imec, Belgium), Sushil Sakhare (imec, Belgium), Peter Debacker (imec, Belgium), Gouri Kar (imec, Belgium), Arnaud Furnemont (imec, Belgium), Nicolas Bueno (Universidad Complutense de Madrid, Spain), José Ignacio Gómez-Pérez (Universidad Complutense de Madrid, Spain), and Christian Tenllado (Universidad Complutense de Madrid, Spain)</i>	
NVDL-Cache: Narrow-Width Value Aware Variable Delay Low-Power Data Cache	264
<i>Nezam Rohbani (Institute for Research in Fundamental Sciences (IPM)), Tapas Kumar Maiti (Hiroshima University), Dondee Navarro (Hiroshima University), Mitiko Miura-Mattausch (Hiroshima University), Hans Jürgen Mattausch (Hiroshima University), and Hirotaka Takatsuka (Mie Fujitsu Semiconductor Limited)</i>	
SpecLock: Speculative Lock Forwarding	273
<i>Pooria M. Yaghini (University of California, Irvine), George Michelogiannakis (Lawrence Berkeley National Laboratory), and Paul V. Gratz (Texas A&M University)</i>	
Fine-Grained Management of Thread Blocks for Irregular Applications	283
<i>Jonathan Beaumont (University of Michigan, USA) and Trevor Mudge (University of Michigan, USA)</i>	

Session 5a: IoT Covert Channel Attacks in Air-Gapped Networks

Red Teaming a Multi-Colored Bluetooth Bulb	293
<i>Ryan Vrecenar (Texas A&M University, USA), Michael Hall (Texas A&M University, USA), Joshua Zshiesche (Texas A&M University, USA), Mahesh Naidu (Texas A&M University, USA), Jeyavijayan Rajendran (Texas A&M University, USA), and Stavros Kalafatis (Texas A&M University, USA)</i>	
Bridging The Gap: Data Exfiltration In Highly Secured Environments Using Bluetooth IoTs	297
<i>Eleonore Carpentier (INSA CVL), Corentin Thomasset (Polytechnique Montréal), and Jeremy Briffaut (INSA CVL - LIFO)</i>	
Covert Data Exfiltration Using Light and Power Channels	301
<i>Patrick Cronin (University of Delaware), Charles Gouert (University of Delaware), Dimitris Mouris (University of Delaware), Nektarios Georgios Tsoutsos (University of Delaware), and Chengmo Yang (University of Delaware)</i>	

Session 5B: Miscellaneous topics in Test, Verification, and Security

Post-Model Validation of Victim DRAM Caches	305
<i>Debiprasanna Sahoo (Indian Institute of Technology Bhubaneswar), Shivani Tripathy (Indian Institute of Technology Bhubaneswar), Manoranjan Satpathy (Indian Institute of Technology Bhubaneswar), and Madhu Mutyam (Indian Institute of Technology Madras)</i>	
Revisiting Capacitor-Based Trojan Design	309
<i>Mohammad Mahdi Bidmeshki (The University of Texas at Dallas, USA), Kiruba Sankaran Subramani (The University of Texas at Dallas, USA), and Yiorgos Makris (The University of Texas at Dallas, USA)</i>	
Formal Modeling and Verification of NAND Flash Memory Supporting Advanced Operations	313
<i>Shivani Tripathy (Indian Institute of Technology Bhubaneswar), Debiprasanna Sahoo (Indian Institute of Technology Bhubaneswar), Manoranjan Satpathy (Indian Institute of Technology Bhubaneswar), and Srinivas Pinisetty (Indian Institute of Technology Bhubaneswar)</i>	
SCARF: Detecting Side-Channel Attacks at Real-Time using Low-Level Hardware Features	N/A
<i>Han Wang (University of California, Davis), Hossein Sayadi (California State University, Long Beach), Setareh Rafatirad (George Mason University), and Houman Homayoun (University of California, Davis)</i>	

Session 6a: Hardware and Software Implementations for Efficient Post Quantum Cryptography

SaberX4: High-Throughput Software Implementation of Saber Key Encapsulation Mechanism	321
<i>Sujoy Sinha Roy (School of Computer Science, University of Birmingham)</i>	
Advances and Challenges of Rank Metric Cryptography Implementations	325
<i>Emanuele Bellini (Technology Innovation Institute), Florian Caullery (Technology Innovation Institute), Rusydi Makarim (Technology Innovation Institute), Marc Manzano (Technology Innovation Institute), Chiara Marcolla (Technology Innovation Institute), and Victor Mateu (Technology Innovation Institute)</i>	

Flexible NTT Accelerators for RLWE Lattice-Based Cryptography	329
<i>Hamid Nejatollahi (University of California, Irvine), Rosario Cammarota (Intel AI, Privacy and Security Research), and Nikil Dutt (University of California, Irvine)</i>	
PlaidML-HE: Acceleration of Deep Learning Kernels to Compute on Encrypted Data	333
<i>Huili Chen (Intel AI Privacy and Security Research, USA), Rosario Cammarota (Intel AI Privacy and Security Research, USA), Felipe Valencia (Advanced Learning and Research Institute, Switzerland), and Francesco Regazzoni (Advanced Learning and Research Institute, Switzerland)</i>	
Power, Area, Speed, and Security (PASS) Trade-Offs of NIST PQC Signature Candidates Using a C to ASIC Design Flow	337
<i>Deepraj Soni (New York University, New York, NY, USA), Mohammed Nabeel (New York University - Abu Dhabi, Abu Dhabi, UAE), Kanad Basu (University of Texas at Dallas - Dallas, Texas, USA), and Ramesh Karri (New York University, New York, NY, USA)</i>	

Session 6B: Architecture & Compilers I

TicToc: Enabling Bandwidth-Efficient DRAM Caching for Both Hits and Misses in Hybrid Memory Systems .	341
<i>Vinson Young (Georgia Institute of Technology), Zeshan A. Chishti (Intel), and Moinuddin K. Qureshi (Georgia Institute of Technology)</i>	
Static Function Prefetching for Efficient Code Management on Scratchpad Memory	350
<i>Youngbin Kim (Yonsei University, Korea), Kyoungwoo Lee (Yonsei University, Korea), and Aviral Shrivastava (Arizona State University, USA)</i>	
Low Power Design through Frequency-Optimized Runtime Micro-Architectural Adaptation	359
<i>Jianqi Chen (The University of Texas at Dallas, USA) and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
HiNUMA: NUMA-Aware Data Placement and Migration in Hybrid Memory Systems	367
<i>Zhuohui Duan (Huazhong University of Science and Technology, China), Haikun Liu (Huazhong University of Science and Technology, China), Xiaofei Liao (Huazhong University of Science and Technology, China), Hai Jin (Huazhong University of Science and Technology, China), Wenbin Jiang (Huazhong University of Science and Technology, China), and Yu Zhang (Huazhong University of Science and Technology, China)</i>	

Session 7A: Advances in Power, Thermal and Timing Aware Optimization

FPGA Energy Efficiency by Leveraging Thermal Margin	376
<i>Behnam Khaleghi (University of California San Diego), Sahand Salamat (University of California San Diego), Mohsen Imani (University of California San Diego), and Tajana Rosing (University of California San Diego)</i>	

A Memetic Algorithm Based PVT Variation-Aware Robust Transistor Sizing Scheme for Power-Delay Optimal Digital Standard Cell Design	385
<i>Mohammed Salman Ahmed (International Institute of Information Technology (IIIT) Hyderabad, India) and Zia Abbas (International Institute of Information Technology (IIIT) Hyderabad, India)</i>	
CSM-NN: Current Source Model Based Logic Circuit Simulation - A Neural Network Approach	393
<i>Mohammad Saeed Abrishami (University of Southern California), Massoud Pedram (University of Southern California), and Shahin Nazarian (University of Southern California)</i>	
Exploiting the Benefits of High-Level Synthesis for Thermal-Aware VLSI Design	401
<i>Jianqi Chen (The University of Texas at Dallas, USA) and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
Efficient Linear System Solution Techniques in the Simulation of Large Dense Mutually Inductive Circuits	405
<i>Charalampos Antoniadis (University of Thessaly, Greece), Milan Mihajlovic (The University of Manchester, United Kingdom), Nestor Evmorfopoulos (University of Thessaly, Greece), Georgios Stamoulis (University of Thessaly, Greece), and Vasilis F. Pavlidis (The University of Manchester, United Kingdom)</i>	
System-Level Optimization of Network-on-Chips for Heterogeneous 3D System-on-Chips	409
<i>Jan Moritz Joseph (Otto-von-Guericke-Universität Magdeburg, Germany), Dominik Ermel (Otto-von-Guericke-Universität Magdeburg, Germany), Lennart Bamberg (Universität Bremen, Germany), Alberto García Oritz (Universität Bremen, Germany), and Thilo Pionteck (Otto-von-Guericke-Universität Magdeburg, Germany)</i>	

Session 7B: EDA for Non-Logic Issues and Non-CMOS Technologies

Constraint-Programmed Initial Sizing of Analog Operational Amplifiers	413
<i>Inga Abel (Technical University of Munich), Maximilian Neuner (Technical University of Munich), and Helmut Graeb (Technical University of Munich)</i>	
Forecast-Based Sample Preparation Algorithm for Unbalanced Splitting Correction on DMFBs	422
<i>Ling-Yen Song (National Chiao Tung University), Yi-Ling Chen (National Chiao Tung University), Yung-Chun Lei (National Chiao Tung University), and Juinn-Dar Huang (National Chiao Tung University)</i>	
A Buffer and Splitter Insertion Framework for Adiabatic Quantum-Flux-Parametron Superconducting Circuits	429
<i>Ruizhe Cai (Northeastern University), Olivia Chen (Yokohama National University), Ao Ren (Northeastern University), Ning Liu (Northeastern University), Nobuyuki Yoshikawa (Yokohama National University), and Yanzhi Wang (Northeastern University)</i>	
PyOCN: A Unified Framework for Modeling, Testing, and Evaluating On-Chip Networks	437
<i>Cheng Tan (Cornell University), Yanghui Ou (Cornell University), Shunning Jiang (Cornell University), Peitian Pan (Cornell University), Christopher Torng (Cornell University), Shady Agwa (Cornell University), and Christopher Batten (Cornell University)</i>	

qCG: A Low-Power Multi-Domain SFQ Logic Design and Verification Framework	446
<i>Shahin Nazarian (University of Southern California, USA), Arash Fayyazi (University of Southern California, USA), and Massoud Pedram (University of Southern California, USA)</i>	

Session 8a: Robust Hardware Design with Machine Learning

A Learning-Based Framework for Automatic Parameterized Verification	450
<i>Yongjian Li (Chinese Academy of Sciences, Institute of Software, The State Key Laboratory of Computer Science, Beijing, China), Jialun Cao (Chinese Academy of Sciences, Institute of Software, The State Key Laboratory of Computer Science, Beijing, China), and Jun Pang (University of Luxembourg, FSTC & SnT, Esch-sur-Alzette, Luxembourg)</i>	
Learning-Based Diversity Estimation: Leveraging the Power of High-Level Synthesis to Mitigate Common-Mode Failure	460
<i>Farah Naz Taher (The University of Texas at Dallas, USA), Anjana Balachandran (The Hong Kong Polytechnic University), and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
Soft Error Resilience in Chip Multiprocessor Cache using a Markov Model Based Re-usability Predictor.....	468
<i>Avishek Choudhury (New Alipore College) and Biplab K Sikdar (IEST Shibpur)</i>	
CHASE: A Customized Time Series Machine Learning Approach for Hardware-Based Stealthy Malware Detection	N/A
<i>Hossein Sayadi (California State University, Long Beach, USA), Yifeng Gao (George Mason University, USA), Hosein Mohammadi Makrani (University of California, Davis, USA), Sai Manoj P D (George Mason University, USA), Avesta Sasan (George Mason University, USA), Jessica Lin (George Mason University, USA), Setareh Rafatirad (George Mason University, USA), and Houman Homayoun (University of California, Davis, USA)</i>	

Session 8B: Accelerators and Machine Learning II

ReNEW: Enhancing Lifetime for ReRAM Crossbar Based Neural Network Accelerators	487
<i>Wen Wen (University of Pittsburgh), Youtao Zhang (University of Pittsburgh), and Jun Yang (University of Pittsburgh)</i>	
Ebird: Elastic Batch for Improving Responsiveness and Throughput of Deep Learning Services	497
<i>Weihao Cui (Shanghai Jiao Tong University, China), Mengze Wei (Shanghai Jiao Tong University, China), Quan Chen (Shanghai Jiao Tong University, China), Xiaoxin Tang (Shanghai University of Finance and Economics, China), Jingwen Leng (Shanghai Jiao Tong University, China), Li Li (Shanghai Jiao Tong University, China), and Mingyi Guo (Shanghai Jiao Tong University, China)</i>	

When Deep Learning Meets the Edge: Auto-Masking Deep Neural Networks for Efficient Machine Learning on Edge Devices	506
<i>Ning Lin (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Hang Lu (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Xing Hu (University of California, Santa Barbara), Jingliang Gao (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Mingzhe Zhang (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), and Xiaowei Li (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences)</i>	
A High-Performance Processing-in-Memory Accelerator for Inline Data Deduplication	515
<i>Young Seo Lee (Korea University, South Korea), Kyung Min Kim (Korea University, South Korea), Ji Heon Lee (Korea University, South Korea), Jeong Hwan Choi (Korea University, South Korea), and Sung Woo Chung (Korea University, South Korea)</i>	

Session 9A: Innovative Circuits for Improved Power Consumption and Performance Characteristics

Low Power Design of Runtime Reconfigurable FPGAs through Contexts Approximations	524
<i>Siyuan Xu (The University of Texas at Dallas, USA) and Benjamin Carrion Schafer (The University of Texas at Dallas, USA)</i>	
Applying Swarm Intelligence to Distributed On-Chip Power Management	532
<i>Divya Pathak (Drexel University) and Ioannis Savidis (Drexel University)</i>	
FeMAT: Exploring In-Memory Processing in Multifunctional FeFET-Based Memory Array	541
<i>Xiaoyu Zhang (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Xiaoming Chen (Institute of Computing Technology, Chinese Academy of Sciences), and Yinhe Han (Institute of Computing Technology, Chinese Academy of Sciences)</i>	
Threshold Logic in a Flash	550
<i>Ankit Wagle (Arizona State University), Gian Singh (Arizona State University), Jinghua Yang (Arizona State University), Sunil Khatri (Texas A&M University), and Sarma Vrudhula (Arizona State University)</i>	

Session 9B: Hardware Security

Adaptive Masking: a Dynamic Trade-off between Energy Consumption and Hardware Security	559
<i>Maxime Montoya (CEA Leti), Thomas Hiscock (CEA Leti), Simone Bacles-Min (CEA Leti), Anca Molnos (CEA Leti), and Jacques Fournier (CEA Leti)</i>	
Cyclic Beneš Network Based Logic Encryption for Mitigating SAT-Based Attacks	567
<i>Saranyu Chattopadhyay (Stanford University) and Rajat Subhra Chakraborty (Indian Institute of Technology Kharagpur)</i>	

PUF-RLA: A PUF-Based Reliable and Lightweight Authentication Protocol Employing Binary String Shuffling	576
<i>Mahmood Azhar Qureshi (Kansas State University) and Arslan Munir (Kansas State University)</i>	
AdapTimer: Hardware/Software Collaborative Timer Resistant to Flush-Based Cache Attacks on ARM-FPGA Embedded SoC	585
<i>Jingquan Ge (University of Chinese Academy of Sciences, China), Neng Gao (Chinese Academy of Sciences, China), Chenyang Tu (Chinese Academy of Sciences, China), Ji Xiang (Chinese Academy of Sciences, China), and Zeyi Liu (Chinese Academy of Sciences, China)</i>	

Session 10A: Miscellaneous Topics in Computer Systems I

Architectural and Cost Implications of the 5G Edge NFV Systems	594
<i>Yang Hu (The University of Texas at Dallas, USA) and Jianda Wang (The University of Texas at Dallas, USA)</i>	
Imbalance-Aware Scheduler for Fast and Secure Ring ORAM Data Retrieval	604
<i>Yuezhi Che (illinois Institute of Technology), Yuan Hong (illinois Institute of Technology), and Rujia Wang (illinois Institute of Technology)</i>	
SPA-SSD: Exploit Heterogeneity and Parallelism of 3D SLC-TLC Hybrid SSD to Improve Write Performance	613
<i>Wenhui Zhang (Huazhong University of Science and Technology), Qiang Cao (Huazhong University of Science and Technology), Hong Jiang (University of Texas at Arlington), Jie Yao (Huazhong University of Science and Technology), Yuanyuan Dong (Alibaba Group.), and Puyuan Yang (Alibaba Group.)</i>	
Mitigating Application Diversity for Allocating a Unified ACC-Rich Platform	622
<i>Jinghan Zhang (Northeastern University, USA), Hamed Tabkhi (University of North Carolina at Charlotte, USA), and Gunar Schirner (Northeastern University, USA)</i>	
VNet: A Versatile Network for Efficient Real-Time Semantic Segmentation	626
<i>Ning Lin (nstitute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Hang Lu (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Jingliang Gao (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences), Shunjie Qiao (University of Hong Kong), and Xiaowei Li (Institute of Computing Technology, Chinese Academy of Sciences; University of Chinese Academy of Sciences)</i>	

Session 10B: Machine Learning Techniques for Innovative Energy-Efficient Solutions

Energy Prediction for Cache Tuning in Embedded Systems	630
<i>Ruben Vazquez (University of Florida), Ann Gordon-Ross (University of Florida), and Greg Stitt (University of Florida)</i>	

Power Management of Wireless Sensor Nodes with Coordinated Distributed Reinforcement Learning	638
<i>Shaswot Shresthamali (The University of Tokyo, Japan), Masaaki Kondo (The University of Tokyo, Japan), and Hiroshi Nakamura (The University of Tokyo, Japan)</i>	
Dynamic Optimization of Battery Health in IoT Networks	648
<i>Kazim Ergun (University of California, San Diego), Raid Ayoub (Intel Corporation), Pietro Mercati (Intel Corporation), and Tajana Rosing (University of California, San Diego)</i>	
Reduced Precision for Energy Efficient FPGA	N/A
<i>Osman Elgawi (University of California, Los Angeles (UCLA), USA) and A. M. Mutawa (Kuwait University, Kuwait)</i>	
BRASIL: A High-Integrity GPGPU Toolchain for Automotive Systems	660
<i>Matina Maria Trompouki (Universitat Politècnica de Catalunya) and Leonidas Kosmidis (Barcelona Supercomputing Center (BSC) and Universitat Politècnica de Catalunya)</i>	

Session 11A: Miscellaneous Topics in Computer Systems II

Balancing Performance and Energy Efficiency of ONoC by Using Adaptive Bandwidth	664
<i>Mingzhe Zhang (State Key Laboratory of Computer Architecture, ICT, CAS, Beijing, China), Lunkai Zhang (Intel Corporation), Frederic T. Chong (University of Chicago), and Zhiyong Liu (State Key Laboratory of Computer Architecture, ICT, CAS, Beijing, China)</i>	
VaLLR: Threshold Voltage Distribution Aware LLR Optimization to Improve LDPC Decoding Performance for 3D TLC NAND Flash	668
<i>Lanlan Cui (Huazhong University of Science and Technology, Wuhan), Fei Wu (Huazhong University of Science and Technology, Wuhan), Xiaojian Liu (DERA Co., Ltd, Shanghai, China), Meng Zhang (Huazhong University of Science and Technology, Wuhan), and Changsheng Xie (Huazhong University of Science and Technology, Wuhan)</i>	
Reinforce Memory Error Protection by Breaking DRAM Disturbance Correlation Within ECC Words	672
<i>Yicheng Wang (University of Illinois at Chicago, USA), Yang Liu (University of Illinois at Chicago, USA), Peiyun Wu (University of Illinois at Chicago, USA), and Zhao Zhang (University of Illinois at Chicago, USA)</i>	
Archivist: A Machine Learning Assisted Data Placement Mechanism for Hybrid Storage Systems	676
<i>Jinting Ren (Chongqing University, China), Xianzhang Chen (Chongqing University, China), Yujuan Tan (Chongqing University, China), Duo Liu (Chongqing University, China), Moming Duan (Chongqing University, China), Liang Liang (Chongqing University, China), and Lei Qiao (Beijing Institute of Control Engineering, China)</i>	
Detecting and Predicting Performance Degradation Caused by Impaired Cache Isolation	680
<i>Yi Zhang (Northeastern University, China), Zhanwei Ling (Northeastern University, China), Ran Cui (Northeastern University, China), Mingsong Lv (Northeastern University, China), Nan Guan (The Hong Kong Polytechnic University, Hong Kong), and Qingxu Deng (Northeastern University, China)</i>	

Session 11B: Processor and Memory Architectures

A Case for Software-Based Adaptive Routing in NUMA Systems	684
<i>Wonjun Song (Electronics and Telecommunications Research Institute, South Korea) and John Kim (Korea Advanced Institute of Science and Technology, South Korea)</i>	
Value Speculation through Equality Prediction	694
<i>Kleovoulos Kalaitzidis (Univ Rennes, Inria, CNRS, IRISA) and André Sez nec (Univ Rennes, Inria, CNRS, IRISA)</i>	
A Novel Convolutional Neural Network Accelerator That Enables Fully-Pipelined Execution of Layers	698
<i>Donghyun Kang (Seoul National University, Republic of Korea), Jintaek Kang (Seoul National University, Republic of Korea), Hyungdal Kwon (Samsung Electronics, Republic of Korea), Hyunsik Park (Samsung Electronics, Republic of Korea), and Soonhoi Ha (Seoul National University, Republic of Korea)</i>	
Freeflow Core: Enhancing Performance of In-Order Cores with Energy Efficiency	702
<i>Raj Kumar Choudhary (Indian Institute of Technology Bombay), Newton Singh (Indian Institute of Technology Bombay), Harideep Nair (Indian Institute of Technology Bombay), Rishabh Rawat (Indian Institute of Technology Bombay), and Virendra Singh (Indian Institute of Technology Bombay)</i>	
Author Index	707