

2023 60th ACM/IEEE Design Automation Conference (DAC 2023)

**San Francisco, California, USA
9-13 July 2023**

Pages 1-858



**IEEE Catalog Number: CFP23DAC-POD
ISBN: 979-8-3503-2349-8**

**Copyright © 2023 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:	CFP23DAC-POD
ISBN (Print-On-Demand):	979-8-3503-2349-8
ISBN (Online):	979-8-3503-2348-1

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

TABLE OF CONTENTS

Map-And-Conquer: Energy-Efficient Mapping of Dynamic Neural Nets onto Heterogeneous MPSoCs.....	1
<i>Halima Bouzidi, Mohanad Odema, Hamza Ouarnoughi, Smail Niar, Mohammad Abdullah Al Faruque</i>	
Specialization Meets Flexibility: A Heterogeneous Architecture for High-Efficiency, High-flexibility AR/VR Processing	7
<i>Arpan Suravi Prasad, Luca Benini, Francesco Conti</i>	
Fine-Grained QoS Control Via Tightly-Coupled Bandwidth Monitoring and Regulation for FPGA-based Heterogeneous SoCs.....	13
<i>G. Brilli, G. Valente, A. Capotondi, P. Burgio, T. Di Masciov, P. Valente, A. Marongiu</i>	
Architecting Selective Refresh Based Multi-Retention Cache for Heterogeneous System (ARMOUR).....	19
<i>Sukarn Agarwal, Shounak Chakraborty, Magnus Sjölander</i>	
Equality Saturation for Datapath Synthesis: A Pathway to Pareto Optimality	25
<i>Ecenur Ustun, Cunxi Yu, Zhiru Zhang</i>	
Rethinking AIG Resynthesis in Parallel	27
<i>Tianji Liu, Evangeline F. Y. Young</i>	
A Database Dependent Framework for K-Input Maximum Fanout-Free Window Rewriting.....	33
<i>Xuliang Zhu, Ruofei Tang, Lei Chen, Xing Li, Xin Huang, Mingxuan Yuan, Weihua Sheng, Jianliang Xu</i>	
Improving Standard-Cell Design Flow Using Factored Form Optimization.....	39
<i>Alessandro Tempia Calvino, Alan Mishchenko, Herman Schmit, Ethan Mahintorabi, Giovanni De Micheli, Xiaoqing Xu</i>	
Don't-Care Aware ESOP Extraction Via Reduced Decomposition-Tree Exploration	45
<i>Chun-Yu Wei, Jie-Hong R. Jiang</i>	
System-Level Innovation for the Next Decade AI Performance, Power, Area with Co-optimization.....	51
<i>Marcus Pan</i>	
PowerPruning: Selecting Weights and Activations for Power-Efficient Neural Network Acceleration.....	53
<i>Richard Petri, Grace Li Zhang, Yiran Chen, Ulf Schlichtmann, Bing Li</i>	
CPE: An Energy-Efficient Edge-Device Training with Multi-dimensional Compression Mechanism	59
<i>Zhou Wang, Jingchuan Wei, Boxiao Han, Hongjun He, Leibo Liu, Shaojun Wei, Shouyi Yin</i>	
Glass Interposer Integration of Logic and Memory Chipllets: PPA and Power/Signal Integrity Benefits.....	65
<i>Pruek Vanna-Iampikul, Lingjun Zhu, Serhat Erdogan, Mohanalingam Kathaperumal, Ravi Agarwal, Ram Gupta, Kevin Rinebold, Sung Kyu Lim</i>	
RoSE: Robust Analog Circuit Parameter Optimization with Sampling-Efficient Reinforcement Learning	71
<i>Jian Gao, Weidong Cao, Xuan Zhang</i>	

Lightning Talk: Model, Framework and Integration for In-Storage Computing with Computational SSDs	77
<i>Tianyu Wang, Jin Xue, Zelin Du, Zhiqi Wang, Yaotian Cui, Zili Shao</i>	
Learning-Based Data Separation for Write Amplification Reduction in Solid State Drives	79
<i>Penghao Sun, Litong You, Shengan Zheng, Wanru Zhang, Ruoyan Ma, Jie Yang, Guanzhong Wang, Feng Zhu, Shu Li, Linpeng Huang</i>	
Reinforcement Learning-Assisted Management for Convertible SSDs	85
<i>Qian Wei, Yi Li, Zhiping Jia, Mengying Zhao, Zhaoyan Shen, Bingzhe Li</i>	
Fair Will Go On: A Collaboration-Aware Fairness Scheme for NVMe SSD in Cloud Storage System	91
<i>Yang Zhou, Fang Wang, Zhan Shi, Dan Feng, Yu Du</i>	
Optimizing the Performance of NDP Operations by Retrieving File Semantics in Storage.....	97
<i>Lin Li, Xianzhang Chen, Jiali Li, Jiapin Wang, Duo Liu, Yujuan Tan, Ao Ren</i>	
Lightning Talk: The Incredible Shrinking Black Box Model	103
<i>Patrick Schaumont</i>	
Hybrid Obfuscation of Chiplet-Based Systems.....	105
<i>Yousef Safari, Pooya Aghanoury, Subramanian S. Iyer, Nader Sehatbakhsh, Boris Vaisband</i>	
ActiWate: Adaptive and Design-Agnostic Active Watermarking for IP Ownership in Modern SoCs.....	111
<i>Zahin Ibnat, M Sadadur Rahman, Mridha Mashahedur Rahman, Hadi Mardani Kamali, Mark Tehranipoor, Farimah Farahmandi</i>	
On the Unpredictability of SPICE Simulations for Side-Channel Leakage Verification of Masked Cryptographic Circuits	117
<i>Kazuki Monta, Makoto Nagata, Josep Balasch, Ingrid Verbauwhede</i>	
An FPGA-Compatible TRNG with Ultra-High Throughput and Energy Efficiency.....	123
<i>Zhaojun Lu, Houjia Qidiao, Qidong Chen, Zhenglin Liu, Jiliang Zhang</i>	
Lightning Talk: Bridging Neuro-Dynamics and Cognition	129
<i>Mohsen Imani</i>	
Neurogenesis Dynamics-Inspired Spiking Neural Network Training Acceleration	131
<i>Shaoyi Huang, Haowen Fang, Kaleel Mahmood, Bowen Lei, Nuo Xu, Bin Lei, Yue Sun, Dongkuan Xu, Wujie Wen, Caiwen Ding</i>	
An Extension to Basis-Hypervectors for Learning from Circular Data in Hyperdimensional Computing.....	137
<i>Igor Nunes, Mike Heddes, Tony Givargis, Alexandru Nicolau</i>	
DistHD: A Learner-Aware Dynamic Encoding Method for Hyperdimensional Classification	143
<i>Junyao Wang, Sitao Huang, Mohsen Imani</i>	
Ising-CF: A Pathbreaking Collaborative Filtering Method Through Efficient Ising Machine Learning	149
<i>Zhuo Liu, Yunan Yang, Zhenyu Pan, Anshujit Sharma, Amit Hasan, Caiwen Ding, Ang Li, Michael Huang, Tong Geng</i>	
Lightning Talk: Scaling Up Quantum Compilation – Challenges and Opportunities.....	155
<i>Jason Cong</i>	

Battle Against Fluctuating Quantum Noise: Compression-Aided Framework to Enable Robust Quantum Neural Network	157
<i>Zhirui Hu, Youzuo Lin, Qiang Guan, Weiwen Jiang</i>	
Compiler Optimization for Quantum Computing Using Reinforcement Learning	163
<i>Nils Quetschlich, Lukas Burgholzer, Robert Wille</i>	
Orchestrating Measurement-Based Quantum Computation Over Photonic Quantum Processors	169
<i>Yingheng Li, Aditya Pawar, Mohadeseh Azari, Yanan Guo, Youtao Zhang, Jun Yang, Kaushik Parasuram Seshadreesan, Xulong Tang</i>	
Scalable Optimal Layout Synthesis for NISQ Quantum Processors	175
<i>Wan-Hsuan Lin, Jason Kimko, Bochen Tan, Nikolaj Bjørner, Jason Cong</i>	
Profile-Driven Banded Smith-Waterman Acceleration for Short Read Alignment	181
<i>Konstantina Koliogeorgi, Dimitrios Soudris, Sotirios Xydis</i>	
Fast FPGA Accelerator of Graph Cut Algorithm with Out-Of-Order Parallel Execution in Folding Grid Architecture	187
<i>Guangyao Yan, Xinzhe Liu, Hui Wang, Yajun Ha</i>	
Optimizing Data Reuse for CGRA Mapping Using Polyhedral-Based Loop Transformations.....	193
<i>Liao Huang, Dajiang Liu</i>	
HAWEN: Hardware Accelerator for Thread Wake-Ups in Linux Event Notification	199
<i>Lars Nolte, Tim Twardzik, Camille Jalier, Zhigang Huang, Jiyuan Shi, Clara Kowalsky, Thomas Wild, Andreas Herkersdorf</i>	
Lightning Talk: All Routes to Timing Closure	205
<i>Iris Hui-Ru Jiang</i>	
RL-CCD: Concurrent Clock and Data Optimization Using Attention-Based Self-Supervised Reinforcement Learning	207
<i>Yi-Chen Lu, Wei-Ting Chan, Deyuan Guo, Sudipto Kundu, Vishal Khandelwal, Sung Kyu Lim</i>	
Restructure-Tolerant Timing Prediction Via Multimodal Fusion	213
<i>Ziyi Wang, Siting Liu, Yuan Pu, Song Chen, Tsung-Yi Ho, Bei Yu</i>	
Critical Paths Prediction Under Multiple Corners Based on BiLSTM Network	219
<i>Qianqian Song, Xu Cheng, Peng Cao</i>	
TOTAL: Multi-Corners Timing Optimization Based on Transfer and Active Learning.....	225
<i>Wei W. Xing, Zheng Xing, Rongqi Lu, Zhelong Wang, Ning Xu, Yuanqing Cheng, Weisheng Zhao</i>	
Ever More Optimized Simulations of Fermionic Systems on a Quantum Computer.....	231
<i>Qingfeng Wang, Ze-Pei Cian, Ming Li, Igor L. Markov, Yunseong Nam</i>	
Optimal Synthesis of Multi-Controlled Qudit Gates	237
<i>Wei Zi, Qian Li, Xiaoming Sun</i>	
Quixote: Improving Fidelity of Quantum Program by Independent Execution of Controlled Gates	243
<i>Enhyeok Jang, Seungwoo Choi, Won Woo Ro</i>	
Hybrid Gate-Pulse Model for Variational Quantum Algorithms	249
<i>Zhiding Liang, Zhixin Song, Jinglei Cheng, Zichang He, Ji Liu, Hanrui Wang, Ruiyang Qin, Yiru Wang, Song Han, Xuehai Qian, Yiyu Shi</i>	

Advances and Trends on On-Chip Compute-in-Memory Macros and Accelerators	255
<i>Jae-Sun Seo</i>	
Processing-In-Hierarchical-Memory Architecture for Billion-Scale Approximate Nearest Neighbor Search.....	257
<i>Zhenhua Zhu, Jun Liu, Guohao Dai, Shulin Zeng, Bing Li, Huazhong Yang, Yu Wang</i>	
A Digital 3D TCAM Accelerator for the Inference Phase of Random Forest	263
<i>Chieh-Lin Tsai, Chun-Feng Wu, Yuan-Hao Chang, Han-Wen Hu, Yung-Chun Lee, Hsiang-Pang Li, Tei-Wei Kuo</i>	
A Convolution Neural Network Accelerator Design with Weight Mapping and Pipeline Optimization.....	269
<i>Lixia Han, Peng Huang, Zheng Zhou, Yiyang Chen, Xiaoyan Liu, Jinfeng Kang</i>	
HAIMA: A Hybrid SRAM and DRAM Accelerator-In-Memory Architecture for Transformer.....	275
<i>Yan Ding, Chubo Liu, Mingxing Duan, Wanli Chang, Keqin Li, Kenli Li</i>	
MeG ² : In-Memory Acceleration for Genome Graphs Analysis.....	281
<i>Yu Huang, Long Zheng, Haifeng Liu, Zhuoran Zhou, Dan Chen, Pengcheng Yao, Qinggang Wang, Xiaofei Liao, Hai Jin</i>	
AmgR: Algebraic Multigrid Accelerated on ReRAM	287
<i>Mingjia Fan, Xiaotian Tian, Yintao He, Junxian Li, Yiru Duan, Xiaozhe Hu, Ying Wang, Zhou Jin, Weifeng Liu</i>	
BP-NTT: Fast and Compact in-SRAM Number Theoretic Transform with Bit-Parallel Modular Multiplication.....	293
<i>Jingyao Zhang, Mohsen Imani, Elaheh Sadredini</i>	
NTT-PIM: Row-Centric Architecture and Mapping for Efficient Number-Theoretic Transform on PIM.....	299
<i>Jaewoo Park, Sugil Lee, Jongeun Lee</i>	
Lightning Talk: Efficient Embedded Machine Learning Deployment on Edge and IoT Devices	305
<i>Sudeep Pasricha</i>	
OmniBoost: Boosting Throughput of Heterogeneous Embedded Devices Under Multi-DNN Workload	307
<i>Andreas Karatzas, Iraklis Anagnostopoulos</i>	
DmTSAR-ILP: Allocating a Unified Domain Platform for Streaming Applications	313
<i>Bruno Morais, Gunar Schirner</i>	
ADAPTIVE: Agent-Based Learning for Bounding Time in Mixed-Criticality Systems.....	319
<i>Behnaz Ranjbar, Ali Hosseinghorban, Akash Kumar</i>	
Response Time Analysis and Optimization of DAG Tasks Exploiting Mutually Exclusive Execution	325
<i>Haochun Liang, Xu Jiang, Nan Guan, Qingqiang He, Wang Yi</i>	
Lightning Talk: Private and Secure Edge AI with Hyperdimensional Computing	331
<i>Xiaofan Yu, Minxuan Zhou, Fatemeh Asgarinejad, Onat Gungor, Baris Aksanli, Tajana Rosing</i>	
PASNet: Polynomial Architecture Search Framework for Two-Party Computation-based Secure Neural Network Deployment.....	333
<i>Hongwu Peng, Shanglin Zhou, Yukui Luo, Nuo Xu, Shijin Duan, Ran Ran, Jiahui Zhao, Chenghong Wang, Tong Geng, Wujie Wen, Xiaolin Xu, Caiwen Ding</i>	

IP Protection in TinyML.....	339
<i>Jinwen Wang, Yuhao Wu, Han Liu, Bo Yuan, Roger Chamberlain, Ning Zhang</i>	
C ² PI: An Efficient Crypto-Clear Two-Party Neural Network Private Inference	345
<i>Yuke Zhang, Dake Chen, Souvik Kundu, Haomei Liu, Ruiheng Peng, Peter A. Beerel</i>	
Primer: Fast Private Transformer Inference on Encrypted Data.....	351
<i>Mengxin Zheng, Qian Lou, Lei Jiang</i>	
ASMCap: An Approximate String Matching Accelerator for Genome Sequence Analysis Based on Capacitive Content Addressable Memory	357
<i>Hongtao Zhong, Zhonghao Chen, Wenqin Huangfu, Chen Wang, Yixin Xu, Tianyi Wang, Yao Yu, Yongpan Liu, Vijaykrishnan Narayanan, Huazhong Yang, Xueqing Li</i>	
Victor: A Variation-Resilient Approach Using Cell-Clustered Charge-domain Computing for High-density High-throughput MLC CiM.....	363
<i>Mingyen Lee, Wenjun Tang, Yiming Chen, Juejian Wu, Hongtao Zhong, Yixin Xu, Yongpan Liu, Huazhong Yang, Vijaykrishnan Narayanan, Xueqing Li</i>	
FSPA: An FeFET-Based Sparse Matrix-Dense Vector Multiplication Accelerator	369
<i>Xiaoyu Zhang, Zerun Li, Rui Liu, Xiaoming Chen, Yinhe Han</i>	
AutoDCIM: An Automated Digital CIM Compiler.....	375
<i>Jia Chen, Fengbin Tu, Kunming Shao, Fengshi Tian, Xiao Huo, Chi-Ying Tsui, Kwang-Ting Cheng</i>	
DSPIMM: A Fully Digital SParse In-Memory Matrix Vector Multiplier for Communication Applications.....	381
<i>Amitesh Sridharan, Fan Zhang, Yang Sui, Bo Yuan, Deliang Fan</i>	
UpPipe: A Novel Pipeline Management on In-Memory Processors for RNA-seq Quantification	387
<i>Liang-Chi Chen, Chien-Chung Ho, Yuan-Hao Chang</i>	
Processor Vulnerability Discovery	393
<i>Yongqiang Lyu, Rihui Sun, Gang Qu</i>	
DriverJar: Lightweight Device Driver Isolation for ARM	396
<i>Huamao Wu, Yuan Chen, Yajin Zhou, Yifei Wang, Lubo Zhang</i>	
EP-ORAM: Efficient NVM-Friendly Path Eviction for Ring ORAM in Hybrid Memory	402
<i>Mehrnoosh Raoufi, Jun Yang, Xulong Tang, Youtao Zhang</i>	
PIMA-LPN: Processing-in-memory Acceleration for Efficient LPN-based Post-Quantum Cryptography.....	408
<i>Lin Ding, Song Bian, Jiliang Zhang</i>	
Mkeycutter: A High-Throughput Key Generator of Classic McEliece on Hardware	414
<i>Yihong Zhu, Wenping Zhu, Chen Chen, Min Zhu, Zhengdong Li, Shaojun Wei, Leibo Liu</i>	
CHAM: A Customized Homomorphic Encryption Accelerator for Fast Matrix-Vector Product	420
<i>Xuanle Ren, Zhaohui Chen, Zhen Gu, Yanheng Lu, Ruiguang Zhong, Wen-Jie Lu, Jiansong Zhang, Yichi Zhang, Hanghang Wu, Xiaofu Zheng, Heng Liu, Tingqiang Chu, Cheng Hong, Changzheng Wei, Dimin Niu, Yuan Xie</i>	
AccShield: A New Trusted Execution Environment with Machine-Learning Accelerators	426
<i>Wei Ren, William Kozlowski, Sandhya Koteswara, Mengmei Ye, Hubertus Franke, Deming Chen</i>	

Lightning Talk: Efficiency and Programmability of DNN Accelerators and GPUs	432
<i>Won Woo Ro</i>	
FIONA: Fine-Grained Incoherent Optical DNN Accelerator Search for Superior Efficiency and Robustness.....	434
<i>Mengquan Li, Kenli Li, Mingfeng Lan, Jie Xiong, Zhuo Tang, Weichen Liu</i>	
Accelerating DNN Inference with Heterogeneous Multi-DPU Engines	440
<i>Zelin Du, Wei Zhang, Zimeng Zhou, Zili Shao, Lei Ju</i>	
FLNA: An Energy-Efficient Point Cloud Feature Learning Accelerator with Dataflow Decoupling	446
<i>Dongxu Lyu, Zhenyu Li, Yuzhou Chen, Ningyi Xu, Guanghui He</i>	
An Efficient Accelerator for Point-Based and Voxel-Based Point Cloud Neural Networks.....	452
<i>Xinhao Yang, Tianyu Fu, Guohao Dai, Shulin Zeng, Kai Zhong, Ke Hong, Yu Wang</i>	
AdaS: A Fast and Energy-Efficient CNN Accelerator Exploiting Bit-Sparsity	458
<i>Xiaolong Lin, Gang Li, Zizhao Liu, Yadong Liu, Fan Zhang, Zhuoran Song, Naifeng Jing, Xiaoyao Liang</i>	
DBPS: Dynamic Block Size and Precision Scaling for Efficient DNN Training Supported by RISC-V ISA Extensions	464
<i>Seunghyun Lee, Jeik Choi, Seockhwan Noh, Jahyun Koo, Jaeha Kung</i>	
On EDA-Driven Learning for SAT Solving	470
<i>Min Li, Zhengyuan Shi, Qiuxia Lai, Sadaf Khan, Shaowei Cai, Qiang Xu</i>	
IFHE: Intermediate-Feature Heterogeneity Enhancement for Image Synthesis in Data-Free Knowledge Distillation.....	476
<i>Yi Chen, Ning Liu, Ao Ren, Tao Yang, Duo Liu</i>	
Robust Tickets Can Transfer Better: Drawing More Transferable Subnetworks in Transfer Learning	482
<i>Yonggan Fu, Ye Yuan, Shang Wu, Jiayi Yuan, Yingyan Celine Lin</i>	
PertNAS: Architectural Perturbations for Memory-Efficient Neural Architecture Search.....	488
<i>Afzal Ahmad, Zhiyao Xie, Wei Zhang</i>	
CSQ: Growing Mixed-Precision Quantization Scheme with Bi-level Continuous Sparsification	494
<i>Lirui Xiao, Huanrui Yang, Zhen Dong, Kurt Keutzer, Li Du, Shanghang Zhang</i>	
Emerging Trends in Multi-Accelerator and Distributed System for ML: Devices, Architectures, Tools and Applications	500
<i>Muhammad Shafique</i>	
PATRONoC: Parallel AXI Transport Reducing Overhead for Networks-On-Chip Targeting Multi-Accelerator DNN Platforms at the Edge	502
<i>Vikram Jain, Matheus Cavalcante, Nazareno Bruschi, Michael Rogenmoser, Thomas Benz, Andreas Kurth, Davide Rossi, Luca Benini, Marian Verhelst</i>	
AdaGL: Adaptive Learning for Agile Distributed Training of Gigantic GNNs.....	508
<i>Ruisi Zhang, Mojan Javaheripi, Zahra Ghodsi, Amit Bleiweiss, Farinaz Koushanfar</i>	
ML-CGRA: An Integrated Compilation Framework to Enable Efficient Machine Learning Acceleration on CGRAs	514
<i>Yixuan Luo, Cheng Tan, Nicolas Bohm Agostini, Ang Li, Antonino Tumeo, Nirav Dave, Tong Geng</i>	

MixPipe: Efficient Bidirectional Pipeline Parallelism for Training Large-Scale Models.....	520
<i>Weigang Zhang, Biyu Zhou, Xuehai Tang, Zhaoxing Wang, Songlin Hu</i>	
MARS: Exploiting Multi-Level Parallelism for DNN Workloads on Adaptive Multi-Accelerator Systems.....	526
<i>Guan Shen, Jieru Zhao, Zeke Wang, Zhe Lin, Wenchao Ding, Chentao Wu, Quan Chen, Minyi Guo</i>	
RESPECT: Reinforcement Learning Based Edge Scheduling on Pipelined Coral Edge TPUs	532
<i>Jiaqi Yin, Yingjie Li, Daniel Robinson, Cunxi Yu</i>	
cVTS: A Constrained Voronoi Tree Search Method for High Dimensional Analog Circuit Synthesis.....	538
<i>Aidong Zhao, Xianan Wang, Zixiao Lin, Zhaori Bi, Xudong Li, Changhao Yan, Fan Yang, Li Shang, Dian Zhou, Xuan Zeng</i>	
Reinforcement Learning-Based Analog Circuit Optimizer Using g_m/I_D for Sizing	544
<i>Minjeong Choi, Youngchang Choi, Kyongsu Lee, Seokhyeong Kang</i>	
Automated Design of Complex Analog Circuits with Multiagent Based Reinforcement Learning	550
<i>Jinxin Zhang, Jiarui Bao, Zhangcheng Huang, Xuan Zeng, Ye Lu</i>	
MTL-Designer: An Integrated Flow for Analysis and Synthesis of Microstrip Transmission Line	556
<i>Qipan Wang, Ping Liu, Liguojiang, Mingjie Liu, Yibo Lin, Runsheng Wang, Ru Huang</i>	
FPDsim: A Structural Simulator for Power Grid Analysis of Flat Panel Display.....	562
<i>Chengtao An, Chunqiao Li, Xiangqi Li, Yangfeng Su, Fan Yang, Xuan Zeng</i>	
Accelerating Sparse LU Factorization with Density-Aware Adaptive Matrix Multiplication for Circuit Simulation	568
<i>Tengcheng Wang, Wenhao Li, Haojie Pei, Yuying Sun, Zhou Jin, Weifeng Liu</i>	
A Matching Based Escape Routing Algorithm with Variable Design Rules and Constraints.....	574
<i>Qinghai Liu, Disi Lin, Chuandong Chen, Huan He, Jianli Chen, Yao-Wen Chang</i>	
Disjoint-Path and Golden-Pin Based Irregular PCB Routing with Complex Constraints	576
<i>Qinghai Liu, Qinfei Tang, Jiarui Chen, Chuandong Chen, Ziran Zhu, Huan He, Jianli Chen, Yao-Wen Chang</i>	
Any-Angle Routing for Redistribution Layers in 2.5D IC Packages	582
<i>Min-Hsuan Chung, Je-Wei Chuang, Yao-Wen Chang</i>	
EDGE: Efficient DAG-Based Global Routing Engine.....	588
<i>Jinwei Liu, Evangeline F. Y. Young</i>	
Pathfinding Model and Lagrangian-Based Global Routing.....	594
<i>Pengju Yao, Ping Zhang, Wenxing Zhu</i>	
OpenDRC: An Efficient Open-Source Design Rule Checking Engine with Hierarchical GPU Acceleration.....	600
<i>Zhuolun He, Yihang Zuo, Jiayi Jiang, Haisheng Zheng, Yuzhe Ma, Bei Yu</i>	
Emerging Hardware Technologies and 3D System Integration for Ubiquitous Machine Intelligence	606
<i>Haitong Li</i>	
Intermittent-Aware Neural Network Pruning	608
<i>Chih-Chia Lin, Chia-Yin Liu, Chih-Hsuan Yen, Tei-Wei Kuo, Pi-Cheng Hsiu</i>	

Occamy: Memory-Efficient GPU Compiler for DNN Inference.....	614
<i>Jaeho Lee, Shinnung Jeong, Seungbin Song, Kunwoo Kim, Heelim Choi, Youngsok Kim, Hanjun Kim</i>	
Uint-Packing: Multiply Your DNN Accelerator Performance Via Unsigned Integer DSP Packing	620
<i>Jingwei Zhang, Meng Zhang, Xinye Cao, Guoqing Li</i>	
Condense: A Framework for Device and Frequency Adaptive Neural Network Models on the Edge	626
<i>Yifan Gong, Pu Zhao, Zheng Zhan, Yushu Wu, Chao Wu, Zhenglun Kong, Minghai Qin, Caiwen Ding, Yanzhi Wang</i>	
NetBooster: Empowering Tiny Deep Learning by Standing on the Shoulders of Deep Giants	632
<i>Zhongzhi Yu, Yonggan Fu, Jiayi Yuan, Haoran You, Yingyan Celine Lin</i>	
WinoTrain: Winograd-Aware Training for Accurate Full 8-bit Convolution Acceleration	638
<i>Pierpaolo Mori, Shambhavi Balamuthu Sampath, Lukas Frickenstein, Manoj-Rohit Vemparala, Nael Fasfous, Alexander Frickenstein, Walter Stechele, Claudio Passerone</i>	
Range-Invariant Approximation of Non-Linear Operations for Efficient BERT Fine-Tuning.....	644
<i>Janghyeon Kim, Janghwan Lee, Jungwook Choi, Jeongho Han, Sangheon Lee</i>	
Hardware-Aware Graph Neural Network Automated Design for Edge Computing Platforms	650
<i>Ao Zhou, Jianlei Yang, Yingjie Qi, Yumeng Shi, Tong Qiao, Weisheng Zhao, Chunming Hu</i>	
Lightning Talk: The Next Wave of High-Level Synthesis.....	656
<i>Deming Chen</i>	
Rubick: A Synthesis Framework for Spatial Architectures Via Dataflow Decomposition	659
<i>Zizhang Luo, Liqiang Lu, Size Zheng, Jieming Yin, Jason Cong, Jianwei Yin, Yun Liang</i>	
An Iterative Method for Mapping-Aware Frequency Regulation in Dataflow Circuits	665
<i>Carmine Rizzi, Andrea Guerrieri, Lana Josipovic</i>	
Automating Constraint-Aware Datapath Optimization Using E-Graphs.....	671
<i>Samuel Coward, George A. Constantinides, Theo Drane</i>	
PIM-HLS: An Automatic Hardware Generation Tool for Heterogeneous Processing-In-Memory-based Neural Network Accelerators	677
<i>Yu Zhu, Zhenhua Zhu, Guohao Dai, Fengbin Tu, Hanbo Sun, Kwang-Ting Cheng, Huazhong Yang, Yu Wang</i>	
Sparse Hamming Graph: A Customizable Network-On-Chip Topology	683
<i>Patrick Iff, Maciej Besta, Matheus Cavalcante, Tim Fischer, Luca Benini, Torsten Hoefler</i>	
Chipelets: How Small is Too Small?.....	689
<i>Alexander Graening, Saptadeep Pal, Puneet Gupta</i>	
Toward Parallelism-Optimal Topology Generation for Wavelength-Routed Optical NoC Designs.....	695
<i>Kuan-Cheng Chen, Yan-Lin Chen, Yu-Sheng Lu, Yao-Wen Chang</i>	
HexaMesh: Scaling to Hundreds of Chipelets with an Optimized Chipelet Arrangement.....	701
<i>Patrick Iff, Maciej Besta, Matheus Cavalcante, Tim Fischer, Luca Benini, Torsten Hoefler</i>	
Faster and Stronger Lossless Compression with Optimized Autoregressive Framework.....	707
<i>Yu Mao, Jingzong Li, Yufei Cui, Jason Chun Xue</i>	

Graph Representation Learning for Microarchitecture Design Space Exploration	713
<i>Xiaoling Yi, Jialin Lu, Xiankui Xiong, Dong Xu, Li Shang, Fan Yang</i>	
TD-Magic: from Pictures of Timing Diagrams to Formal Specifications	719
<i>Jie He, Dejan Nickovic, Ezio Bartocci, Radu Grosu</i>	
Stochastic-HMDs: Adversarial-Resilient Hardware Malware Detectors Via Undervolting	725
<i>Md Shohidul Islam, Ihsen Alouani, Khaled N. Khasawneh</i>	
Heterogeneous Reconfigurable Accelerators: Trends and Perspectives	731
<i>Wayne Luk</i>	
HDSuper: Algorithm-Hardware Co-design for Light-weight High-quality Super-Resolution Accelerator	733
<i>Liang Chang, Xin Zhao, Dongqi Fan, Zhicheng Hu, Jun Zhou</i>	
DARIC: A Data Reuse-Friendly CGRA for Parallel Data Access Via Elastic FIFOs.....	739
<i>Dajiang Liu, Di Mou, Rong Zhu, Yan Zhuang, Jiaying Shang, Jiang Zhong, Shouyi Yin</i>	
Orchestrated Scheduling and Partitioning for Improved Address Translation in GPUs.....	745
<i>Bingyao Li, Yueqi Wang, Xulong Tang</i>	
Bumblebee: A MemCache Design for Die-Stacked and Off-chip Heterogeneous Memory Systems	751
<i>Yifan Hua, Shengan Zheng, Ji Yin, Weidong Chen, Linpeng Huang</i>	
Lightning Talk: A Perspective on Neuromorphic Computing	757
<i>Deepika Sharma, Adarsh Kumar Kosta, Kaushik Roy</i>	
ANAS: Asynchronous Neuromorphic Hardware Architecture Search Based on a System-Level Simulator	759
<i>Jian Zhang, Jilin Zhang, Dexuan Huo, Hong Chen</i>	
Input-Aware Dynamic Timestep Spiking Neural Networks for Efficient In-Memory Computing.....	765
<i>Yuhang Li, Abhishek Moitra, Tamar Geller, Priyadarshini Panda</i>	
Neuromorphic Swarm on RRAM Compute-In-Memory Processor for Solving QUBO Problem.....	771
<i>Ashwin Sanjay Lele, Muya Chang, Samuel D. Spetalnick, Brian Crafton, Arijit Raychowdhury, Yan Fang</i>	
ACGraph: Accelerating Streaming Graph Processing Via Dependence Hierarchy	777
<i>Zihan Jiang, Fubing Mao, Yapu Guo, Xu Liu, Haikun Liu, Xiaofei Liao, Hai Jin, Wei Zhang</i>	
MPass: Bypassing Learning-Based Static Malware Detectors	783
<i>Jialai Wang, Wenjie Qu, Yi Rong, Han Qiu, Qi Li, Zongpeng Li, Chao Zhang</i>	
VideoFlip: Adversarial Bit Flips for Reducing Video Service Quality	789
<i>Jung-Woo Chang, Mojan Javaheripi, Farinaz Koushanfar</i>	
HyperAttack: An Efficient Attack Framework for HyperDimensional Computing	795
<i>Fangxin Liu, Haoming Li, Yongbiao Chen, Tao Yang, Li Jiang</i>	
Physics-Informed Optical Kernel Regression Using Complex-valued Neural Fields	801
<i>Guojin Chen, Zehua Pei, Haoyu Yang, Yuzhe Ma, Bei Yu, Martin Wong</i>	
Efficient ILT Via Multi-Level Lithography Simulation.....	807
<i>Shuyuan Sun, Fan Yang, Bei Yu, Li Shang, Xuan Zeng</i>	

Graph-Based Simultaneous Placement and Routing for Two-Dimensional Directed Self-Assembly Technology	813
<i>Wei-Hsu Chen, Yao-Wen Chang</i>	
Lamellar DSA-Aware Detailed Routing Considering Double Patterning and Short Template Minimization	819
<i>Kuei-Lin Wu, Shao-Yun Fang</i>	
Architecture 2.0: Challenges and Opportunities.....	825
<i>Vijay Janapa Reddi, Amir Yazdanbakhsh</i>	
Muffin: A Framework Toward Multi-Dimension AI Fairness by Uniting Off-the-Shelf Models.....	827
<i>Yi Sheng, Junhuan Yang, Lei Yang, Yiyu Shi, Jingtong Hu, Weiwen Jiang</i>	
Fast Adversarial Training with Dynamic Batch-Level Attack Control.....	833
<i>Jaewon Jung, Jaeyong Song, Hongsun Jang, Hyeyoon Lee, Kanghyun Choi, Noseong Park, Jinho Lee</i>	
Unified Agile Accuracy Assessment in Computing-In-Memory Neural Accelerators by Layerwise Dynamical Isometry	839
<i>Xuan-Jun Chen, Cynthia Kuan, Chia-Lin Yang</i>	
Physics-Aware Roughness Optimization for Diffractive Optical Neural Networks.....	845
<i>Shanglin Zhou, Yingjie Li, Minhan Lou, Weilu Gao, Zhijie Shi, Cunxi Yu, Caiwen Ding</i>	
Lightning Talk 21: EDA with ML, Rule-Based, Or Both?	851
<i>Youngsoo Shin</i>	
PSMiner: A Pattern-Aware Accelerator for High-Performance Streaming Graph Pattern Mining.....	853
<i>Hao Qi, Yu Zhang, Ligang He, Kang Luo, Jun Huang, Haoyu Lu, Jin Zhao, Hai Jin</i>	
Instant-NeRF: Instant On-Device Neural Radiance Field Training Via Algorithm-Accelerator Co-Designed Near-Memory Processing	859
<i>Yang Katie Zhao, Shang Wu, Jingqun Zhang, Sixu Li, Chaojian Li, Yingyan Celine Lin</i>	
Bit-Serial Cache: Exploiting Input Bit Vector Repetition to Accelerate Bit-Serial Inference	865
<i>Yun-Chen Lo, Ren-Shuo Liu</i>	
SaGraph: A Similarity-Aware Hardware Accelerator for Temporal Graph Processing	871
<i>Jin Zhao, Yu Zhang, Jian Cheng, Yiyang Wu, Chuyue Ye, Hui Yu, Zhiying Huang, Hai Jin, Xiaofei Liao, Lin Gu, Haikun Liu</i>	
Lightning Talk: Power and Performance Reconciliation – from Tradeoff to Win-Win.....	877
<i>Jiang Hu</i>	
DeepOHeat: Operator Learning-Based Ultra-fast Thermal Simulation in 3D-IC Design	879
<i>Ziyue Liu, Yixing Li, Jing Hu, Xinling Yu, Shinyu Shiau, Xin Ai, Zhiyu Zeng, Zheng Zhang</i>	
STRIVE: Enabling Choke Point Detection and Timing Error Resilience in a Low-Power Tensor Processing Unit.....	885
<i>Noel Daniel Gundi, Zinnia Muntaha Mowri, Andrew Chamberlin, Sanghamitra Roy, Koushik Chakraborty</i>	
PROPHET: Predictive On-Chip Power Meter in Hardware Accelerator for DNN	891
<i>Jian Peng, Tingyuan Liang, Zhiyao Xie, Wei Zhang</i>	

Machine Learning-Based Thermally-Safe Cache Contention Mitigation in Clustered Manycores.....	897
<i>Mohammed Bakr Sikal, Heba Khdr, Martin Rapp, Jörg Henkel</i>	
COSA:Co-Operative Systolic Arrays for Multi-head Attention Mechanism in Neural Network Using Hybrid Data Reuse and Fusion Methodologies.....	903
<i>Zhican Wang, Gang Wang, Honglan Jiang, Ningyi Xu, Guanghui He</i>	
TF-MVP: Novel Sparsity-Aware Transformer Accelerator with Mixed-Length Vector Pruning	909
<i>Eunji Yoo, Gunho Park, Jung Gyu Min, Se Jung Kwon, Baeseong Park, Dongsoo Lee, Youngjoo Lee</i>	
Efficient Transformer Inference with Statically Structured Sparse Attention.....	915
<i>Steve Dai, Hasan Genc, Rangharajan Venkatesan, Brucek Khailany</i>	
Memory and Computation Coordinated Mapping of DNNs onto Complex Heterogeneous SoC	921
<i>Size Zheng, Siyuan Chen, Yun Liang</i>	
APPEND: Rethinking ASIP Synthesis in the Era of AI.....	927
<i>Cangyuan Li, Ying Wang, Huawei Li, Yinhe Han</i>	
Venus: A Versatile Deep Neural Network Accelerator Architecture Design for Multiple Applications	933
<i>Jiaqi Yang, Hao Zheng, Ahmed Louri</i>	
When Monte-Carlo Dropout Meets Multi-Exit: Optimizing Bayesian Neural Networks on FPGA	939
<i>Hongxiang Fan, Mark Chen, Liam Castelli, Zhiqiang Que, He Li, Kenneth Long, Wayne Luk</i>	
Design Automation for Cryogenic CMOS Circuits	945
<i>Victor M. Van Santen, Marcel Walter, Florian Klemme, Shivendra Singh Parihar, Girish Pahwa, Yogesh S. Chauhan, Robert Wille, Hussam Amrouch</i>	
Cryogenic In-Memory Matrix-Vector Multiplication Using Ferroelectric Superconducting Quantum Interference Device (FE-SQUID).....	951
<i>Shamiul Alam, Jack Hutchins, Md Shafayat Hossain, Kai Ni, Vijaykrishnan Narayanan, Ahmedullah Aziz</i>	
Thermal Scaffolding for Ultra-Dense 3D Integrated Circuits	957
<i>Dennis Rich, Anna Kasperovich, Mohamadali Malakoutian, Robert M. Radway, Shiho Hagiwara, Takahide Yoshikawa, Srabanti Chowdhury, Subhasish Mitra</i>	
CorePUM: Efficient Processing Using Cross-Point Memory Via Cooperative Row-Column Access Pipelining and Adaptive Timing Optimization in Subarrays	963
<i>Chengning Wang, Dan Feng, Wei Tong, Jingning Liu</i>	
Compact and High-Performance TCAM Based on Scaled Double-Gate FeFETs.....	969
<i>Liu Liu, Shubham Kumar, Simon Thomann, Hussam Amrouch, Xiaobo Sharon Hu</i>	
A Fast Secure Deletion Strategy for High-Density Flash Memory Through WOM-v Codes.....	975
<i>Jinhua Cui, Kai Tang, Laurence T. Yang</i>	
How to Boost Deep Neural Networks for Computer Vision	981
<i>Soonhoi Ha</i>	
Flex-SFU: Accelerating DNN Activation Functions by Non-Uniform Piecewise Approximation.....	987
<i>Enrico Reggiani, Renzo Andri, Lukas Cavigelli</i>	
A Memory-Efficient Edge Inference Accelerator with XOR-based Model Compression.....	989
<i>Hyunseung Lee, Jihoon Hong, Soosung Kim, Seung Yul Lee, Jae W. Lee</i>	

Dynamic Sparse Training Via Balancing the Exploration-Exploitation Trade-off	995
<i>Shaoyi Huang, Bowen Lei, Dongkuan Xu, Hongwu Peng, Yue Sun, Mimi Xie, Caiwen Ding</i>	
Algorithms and Hardware for Efficient Processing of Logic-Based Neural Networks	1001
<i>Jingkai Hong, Arash Fayyazi, Amirhossein Esmaili, Mahdi Nazemi, Massoud Pedram</i>	
HBP: Hierarchically Balanced Pruning and Accelerator Co-Design for Efficient DNN Inference	1007
<i>Ao Ren, Yuhao Wang, Tao Zhang, Jiaying Shi, Duo Liu, Xianzhang Chen, Yujuan Tan, Yuan Xie</i>	
Efficient Non-Linear Adder for Stochastic Computing with Approximate Spatial-Temporal Sorting Network	1013
<i>Yixuan Hu, Tengyu Zhang, Meng Li, Renjie Wei, Liangzhen Lai, Yuan Wang, Runsheng Wang, Ru Huang</i>	
Don't Cross Me! Cross-Layer System Security	1019
<i>Najmeh Nazari, Chongzhou Fang, Sai Manoj P D, Houman Homayoun</i>	
Stalker: A Framework to Analyze Fragility of Cryptographic Libraries Under Hardware Fault Models	1021
<i>Guorui Xu, Fan Zhang, Xinjie Zhao, Yuan Chen, Shize Guo, Kui Ren</i>	
PTStore: Lightweight Architectural Support for Page Table Isolation	1027
<i>Wende Tan, Yangyu Chen, Yuan Li, Ying Liu, Jianping Wu, Yu Ding, Chao Zhang</i>	
Leaky MDU: ARM Memory Disambiguation Unit Uncovered and Vulnerabilities Exposed	1033
<i>Chang Liu, Yongqiang Lyu, Haixia Wang, Pengfei Qiu, Dapeng Ju, Gang Qu, Dongsheng Wang</i>	
MES-Attacks: Software-Controlled Covert Channels Based on Mutual Exclusion and Synchronization	1039
<i>Chaoqun Shen, Jiliang Zhang, Gang Qu</i>	
ZKROWN: Zero Knowledge Right of Ownership for Neural Networks	1045
<i>Nojan Sheybani, Zahra Ghodsi, Ritvik Kapila, Farinaz Koushanfar</i>	
Smart Detection of Obfuscated Thermal Covert Channel Attacks in Many-Core Processors	1051
<i>Jeferson González-Gómez, Mohammed Bakr Sikal, Heba Khdr, Lars Bauer, Jörg Henkel</i>	
Mitigating Distribution Shift for Congestion Optimization in Global Placement	1057
<i>Su Zheng, Lancheng Zou, Siting Liu, Yibo Lin, Bei Yu, Martin Wong</i>	
PUFFER: A Routability-Driven Placement Framework Via Cell Padding with Multiple Features and Strategy Exploration	1063
<i>Zhijie Cai, Peng Zou, Zhengtao Wu, Xingyu Tong, Jun Yu, Jianli Chen, Yao-Wen Chang</i>	
On a Moreau Envelope Wirelength Model for Analytical Global Placement	1065
<i>Peiyu Liao, Hongduo Liu, Yibo Lin, Bei Yu, Martin Wong</i>	
Global Floorplanning Via Semidefinite Programming	1071
<i>Wei Li, Fangzhou Wang, José M. F. Moura, R. D. Blanton</i>	
Mixed-Cell-height Placement with Minimum-Implant-Area and Drain-to-Drain Abutment Constraints	1077
<i>Guohao Chen, Zheng Zeng, Benchao Zhu, Jiawei Li, Kun Wang, Jun Yu, Jianli Chen</i>	

Concurrent Sign-Off Timing Optimization Via Deep Steiner Points Refinement	1083
<i>Siting Liu, Ziyi Wang, Fangzhou Liu, Yibo Lin, Bei Yu, Martin Wong</i>	
XPert: Peripheral Circuit & Neural Architecture Co-Search for Area and Energy-efficient Xbar-based Computing	1089
<i>Abhishek Moitra, Abhiroop Bhattacharjee, Youngeun Kim, Priyadarshini Panda</i>	
TFix: Exploiting the Natural Redundancy of Ternary Neural Networks for Fault Tolerant In-Memory Vector Matrix Multiplication	1095
<i>Akul Malhotra, Chunguang Wang, Sumeet Kumar Gupta</i>	
Morphable CIM: Improving Operation Intensity and Depthwise Capability for SRAM-CIM Architecture	1101
<i>Yun-Chen Lo, Ren-Shuo Liu</i>	
In-Memory Neural Network Accelerator Based on eDRAM Cell with Enhanced Retention Time	1107
<i>Inhwan Lee, Eunhwan Kim, Nameun Kang, Hyunmyung Oh, Jae-Joon Kim</i>	
ReRAM-Based Graph Attention Network with Node-centric Edge Searching and Hamming Similarity	1113
<i>Ruibin Mao, Xia Sheng, Catherine Graves, Cong Xu, Can Li</i>	
Accelerating Sparse Attention with a Reconfigurable Non-Volatile Processing-In-Memory Architecture	1119
<i>Qilin Zheng, Shiyu Li, Yitu Wang, Zirui Li, Yiran Chen, Hai Helen Li</i>	
Property-Based Timing Analysis and Optimization for Complex Cyber-Physical Real-Time Systems	1125
<i>Jian-Jia Chen, Niklas Ueter, Mario Günzel, Georg Von Der Brüggen, Tei-Wei Kuo</i>	
BlueFace: Integrating an Accelerator into the Core's Pipeline Through Algorithm-Interface Co-Design for Real-Time SoCs	1127
<i>Zhe Jiang, Nathan Fisher, Nan Guan, Zheng Dong</i>	
Holistic WCRT Analysis for Global Fixed-Priority Preemptive Multiprocessor Scheduling	1133
<i>Guoqi Xie, Chenglai Xiong, Wei Wu, Renfa Li, Wanli Chang</i>	
A Universal Method for Task Allocation on FP-FPS Multiprocessor Systems with Spin Locks	1139
<i>Shuai Zhao, Nan Chen, Yinjie Fang, Zhao Li, Wanli Chang</i>	
SPET: Transparent SRAM Allocation and Model Partitioning for Real-Time DNN Tasks on Edge TPU	1145
<i>Changhun Han, Hoon Sung Chwa, Kilho Lee, Sangeun Oh</i>	
Reaction Time Analysis of Event-Triggered Processing Chains with Data Refreshing	1151
<i>Yue Tang, Nan Guan, Xu Jiang, Zheng Dong, Wang Yi</i>	
Fault Tolerance in Time-Sensitive Networking with Mixed-Critical Traffic	1157
<i>Wenhong Ma, Xiangzhen Xiao, Guoqi Xie, Nan Guan, Yu Jiang, Wanli Chang</i>	
NMExplorer: An Efficient Exploration Framework for DIMM-Based Near-Memory Tensor Reduction	1163
<i>Cong Li, Zhe Zhou, Xingchen Li, Guangyu Sun, Dimin Niu</i>	
PIMCOMP: A Universal Compilation Framework for Crossbar-Based PIM DNN Accelerators	1169
<i>Xiaotian Sun, Xinyu Wang, Wanqian Li, Lei Wang, Yinhe Han, Xiaoming Chen</i>	

BWA-NIMC: Budget-based Workload Allocation for Hybrid Near/In-Memory-Computing	1175
<i>Chi-Tse Huang, Cheng-Yang Chang, Yu-Chuan Chuang, An-Yeu Andy Wu</i>	
Rethinking Programming Frameworks for In-Storage Processing	1181
<i>Yu-Chia Liu, Kuan-Chieh Hsu, Hung-Wei Tseng</i>	
ZoneKV: A Space-Efficient Key-Value Store for ZNS SSDs	1187
<i>Mingchen Lu, Peiquan Jin, Xiaoliang Wang, Yongping Luo, Kuankuan Guo</i>	
HUNTER: Releasing Persistent Memory Write Performance with a Novel PM-DRAM Collaboration Architecture	1193
<i>Yanqi Pan, Yifeng Zhang, Wen Xia, Xiangyu Zou, Cai Deng</i>	
Lightning Talk: The New Era of Computational Cognitive Intelligence	1199
<i>Nikil Dutt, Bryan Donyanavard</i>	
On-Device Unsupervised Image Segmentation	1201
<i>Junhuan Yang, Yi Sheng, Yuzhou Zhang, Weiwen Jiang, Lei Yang</i>	
Comprehensive Integration of Hyperdimensional Computing with Deep Learning Towards Neuro- Symbolic AI	1207
<i>Hyunsei Lee, Jiseung Kim, Hanning Chen, Ariela Zeira, Narayan Srinivasa, Mohsen Imani, Yeseong Kim</i>	
PseudoSC: A Binary Approximation to Stochastic Computing Within Latent Operation-Space for Ultra-Lightweight on-Edge DNNs	1213
<i>Jihe Wang, Zhaoqing Wang, Danghui Wang</i>	
Fault-Tolerance-Oriented Physical Design for Fully Programmable Valve Array Biochips	1219
<i>Genggen Liu, Yuhan Zhu, Wenzhong Guo, Xing Huang</i>	
STCG: State-Aware Test Case Generation for Simulink Models	1225
<i>Zhuo Su, Zehong Yu, Dongyan Wang, Yixiao Yang, Rui Wang, Wanli Chang, Aiguo Cui, Yu Jiang</i>	
UpTime: Towards Flow-Based In-Memory Computing with High Fault-Tolerance	1231
<i>Sven Thijssen, Muhammad Rashedul Haq Rashed, Sumit Kumar Jha, Rickard Ewetz</i>	
Fault Injection in Native Logic-In-Memory Computation on Neuromorphic Hardware	1237
<i>Felix Staudigl, Thorben Fetz, Rebecca Pelke, Dominik Sisejkovic, Jan Moritz Joseph, Leticia Bolzani Pöhls, Rainer Leupers</i>	
Seeking the Yield Barrier: High-Dimensional SRAM Evaluation Through Optimal Manifold	1243
<i>Yanfang Liu, Guohao Dai, Wei W. Xing</i>	
Toward Optimal Filler Cell Insertion with Complex Implant Layer Constraints	1249
<i>Peng Zou, Guohao Chen, Zhifeng Lin, Jun Yu, Jianli Chen</i>	
DiffPattern: Layout Pattern Generation Via Discrete Diffusion	1255
<i>Zixiao Wang, Yunheng Shen, Wenqian Zhao, Yang Bai, Guojin Chen, Farzan Farnia, Bei Yu</i>	
LRS DP: Low-Rank SDP for Triple Patterning Lithography Layout Decomposition	1261
<i>Yu Zhang, Yifan Chen, Zhonglin Xie, Hong Xu, Zaiwen Wen, Yibo Lin, Bei Yu</i>	
Layout Decomposition Via Boolean Satisfiability	1267
<i>Hongduo Liu, Peiyu Liao, Mengchuan Zou, Bowen Pang, Xijun Li, Mingxuan Yuan, Tsung-Yi Ho, Bei Yu</i>	

Lightning Talk 6: Bringing Together Foundation Models and Edge Devices	1273
<i>Nick John Eliopoulos, Yung-Hsiang Lu</i>	
SpMMPlu: A Compiler Plug-In with Sparse IR for Efficient Sparse Matrix Multiplication.....	1275
<i>Tao Yang, Yiyuan Zhou, Qidong Tang, Feng Xu, Hui Ma, Jieru Zhao, Li Jiang</i>	
EENet: Energy Efficient Neural Networks with Run-Time Power Management.....	1281
<i>Xiangjie Li, Yingtao Shen, An Zou, Yehan Ma</i>	
EagleRec: Edge-Scale Recommendation System Acceleration with Inter-Stage Parallelism Optimization on GPUs	1287
<i>Yongbo Yu, Fuxun Yu, Xiang Sheng, Chenchen Liu, Xiang Chen</i>	
Energy-Efficient On-Chip Training for Customized Home-based Rehabilitation Systems.....	1293
<i>A. Alper Goksoy, Sizhe An, Umit Y. Ogras</i>	
Lightning Talk: Can Memory Technologies Meet Demands of Data Abundant Applications?	1299
<i>Vijaykrishnan Narayanan</i>	
Lift: Exploiting Hybrid Stacked Memory for Energy-Efficient Processing of Graph Convolutional Networks	1301
<i>Jiaxian Chen, Zhaoyu Zhong, Kaoyi Sun, Chenlin Ma, Rui Mao, Yi Wang</i>	
DRPTM: A Decoupled Read-Efficient High-scalable Persistent Transactional Memory.....	1307
<i>Wenkai Liang, Hao Hu, Xiangyu Zou, Wen Xia, Yanqi Pan</i>	
Correlation-Guided Placement for Nonvolatile FPGAs	1313
<i>Mengying Zhao, Fanjin Xu, Huichuan Zheng, Hao Zhang, Yuqing Xiong, Zhiping Jia, Xiaojun Cai</i>	
RMP-MEM: A HW/SW Reconfigurable Multi-Port Memory Architecture for Multi-PEA Oriented CGRA.....	1319
<i>Qidie Wu, Jiangyuan Gu, Youxu Lin, Boxiao Han, Hongjun He, Yang Hu, Leibo Liu, Shaojun Wei, Shouyi Yin</i>	
TiPU: A Spatial-Locality-Aware Near-Memory Tile Processing Unit for 3D Point Cloud Neural Network.....	1325
<i>Jiapei Zheng, Hao Jiang, Xinkai Nie, Zhangcheng Huang, Chixiao Chen, Qi Liu</i>	
R-TOSS: A Framework for Real-Time Object Detection Using Semi-Structured Pruning	1331
<i>Abhishek Balasubramaniam, Febin Sunny, Sudeep Pasricha</i>	
BLITZCRANK: Factor Graph Accelerator for Motion Planning.....	1337
<i>Yuhui Hao, Yiming Gan, Bo Yu, Qiang Liu, Shao-Shan Liu, Yuhao Zhu</i>	
BERRY: Bit Error Robustness for Energy-Efficient Reinforcement Learning-Based Autonomous Systems.....	1343
<i>Zishen Wan, Nandhini Chandramoorthy, Karthik Swaminathan, Pin-Yu Chen, Vijay Janapa Reddi, Arijit Raychowdhury</i>	
Lightning Talk: Memory-Centric Computing.....	1349
<i>Onur Mutlu</i>	
VPIM: Efficient Virtual Address Translation for Scalable Processing-In-Memory Architectures	1351
<i>Amel Fatima, Sihang Liu, Korakit Seemakhupt, Rachata Ausavarungnirun, Samira Khan</i>	

PUMICE: Processing-Using-Memory Integration with a Scalar Pipeline for Symbiotic Execution.....	1357
<i>Socrates S. Wong, Cecilio C. Tamarit, José F. Martínez</i>	
Hardware Support for Durable Atomic Instructions for Persistent Parallel Programming.....	1363
<i>Khan Shaikhul Hadi, Naveed Ul Mustafa, Mark Heinrich, Yan Solihin</i>	
Sidekick: Near Data Processing for Clustering Enhanced by Automatic Memory Disaggregation	1369
<i>Sanghoon Lee, Jongho Park, Minho Ha, Byung Il Koh, Kyoung Park, Yeseong Kim</i>	
A High-Accurate Multi-objective Exploration Framework for Design Space of CPU.....	1375
<i>Duo Wang, Mingyu Yan, Xin Liu, Mo Zou, Tianyu Liu, Wenming Li, Xiaochun Ye, Dongrui Fan</i>	
A Comprehensive Automated Exploration Framework for Systolic Array Designs.....	1381
<i>Suhail Basalama, Jie Wang, Jason Cong</i>	
High Performance, Low Power Matrix Multiply Design on ACAP: from Architecture, Design Challenges and DSE Perspectives	1387
<i>Jinming Zhuang, Zhuoping Yang, Peipei Zhou</i>	
A Model-Specific End-to-End Design Methodology for Resource-Constrained TinyML Hardware.....	1393
<i>Yanchi Dong, Tianyu Jia, Kaixuan Du, Yiqi Jing, Qijun Wang, Pixian Zhan, Yadong Zhang, Fengyun Yan, Yufei Ma, Yun Liang, Le Ye, Ru Huang</i>	
AccALS: Accelerating Approximate Logic Synthesis by Selection of Multiple Local Approximate Changes	1399
<i>Xuan Wang, Sijun Tao, Jingjing Zhu, Yiyu Shi, Weikang Qian</i>	
Lightweight Structural Choices Operator for Technology Mapping	1405
<i>Antoine Grosnit, Matthieu Zimmer, Rasul Tutunov, Xing Li, Lei Chen, Fan Yang, Mingxuan Yuan, Haitham Bou-Ammar</i>	
AGD: A Learning-Based Optimization Framework for EDA and Its Application to Gate Sizing.....	1411
<i>Phuoc Pham, Jaeyong Chung</i>	
RL-MUL: Multiplier Design Optimization with Deep Reinforcement Learning.....	1417
<i>Dongsheng Zuo, Yikang Ouyang, Yuzhe Ma</i>	
SE3: Sequential Equivalence Checking for Non-Cycle-Accurate Design Transformations †.....	1423
<i>You Li, Guannan Zhao, Yunqi He, Hai Zhou</i>	
Formal Verification of Restoring Dividers Made Fast and Simple.....	1429
<i>Jiteshri Dasari, Maciej Ciesielski</i>	
G-QED: Generalized QED Pre-silicon Verification Beyond Non-Interfering Hardware Accelerators	1435
<i>Saranya Chattopadhyay, Keerthikumara Devarajegowda, Bihan Zhao, Florian Lonsing, Brandon A. D'Agostino, Ioanna Vavelidou, Vijay D. Bhatt, Sebastian Prebeck, Wolfgang Ecker, Caroline Trippel, Clark Barrett, Subhasish Mitra</i>	
Towards a Formally Verified Fully Homomorphic Encryption Compute Engine*	1441
<i>Jeremy Casas, Zhenkun Yang, Wen Wang, Jin Yang, Adwait Godbole</i>	
Lightning Talk: Unlocking the Potential of the Analog Domain: Exploring the Next Frontier in Hardware Security	1447
<i>Waleed Khalil</i>	

ChaosINTC: A Secure Interrupt Management Mechanism Against Interrupt-Based Attacks on TEE	1448
<i>Yifan Zhu, Peinan Li, Lutan Zhao, Dan Meng, Rui Hou</i>	
SCAGuard: Detection and Classification of Cache Side-Channel Attacks Via Attack Behavior Modeling and Similarity Comparison	1454
<i>Limin Wang, Lei Bu, Fu Song</i>	
AVX Timing Side-Channel Attacks Against Address Space Layout Randomization	1460
<i>Hyunwoo Choi, Suryeon Kim, Seungwon Shin</i>	
GDSII-Guard: ECO Anti-Trojan Optimization with Exploratory Timing-Security Trade-Offs	1466
<i>Xinming Wei, Jiayi Zhang, Guojie Luo</i>	
ALMOST: Adversarial Learning to Mitigate Oracle-Less ML Attacks Via Synthesis Tuning	1472
<i>Animesh B. Chowdhury, Lilas Alrahis, Luca Collini, Johann Knechtel, Ramesh Karri, Siddharth Garg, Ozgur Sinanoglu, Benjamin Tan</i>	
SimLL: Similarity-Based Logic Locking Against Machine Learning Attacks	1478
<i>Subhajit Dutta Chowdhury, Kaixin Yang, Pierluigi Nuzzo</i>	
HammerDodger: A Lightweight Defense Framework Against RowHammer Attack on DNNs	1484
<i>Cheng Gongye, Yukui Luo, Xiaolin Xu, Yunsi Fei</i>	
NNTesting: Neural Network Fault Attacks Detection Using Gradient-Based Test Vector Generation	1490
<i>Antian Wang, Bingyin Zhao, Weihang Tan, Yingjie Lao</i>	
ExploreFault: Identifying Exploitable Fault Models in Block Ciphers with Reinforcement Learning	1496
<i>Hao Guo, Sayandeep Saha, Vasudev Gohil, Satwik Patnaik, Debdeep Mukhopadhyay, Jeyavijayan Jv Rajendran</i>	
Power Side-Channel Vulnerability Assessment of Lightweight Cryptographic Scheme, XOODYAK	1502
<i>Anupam Golder, Debayan Das, Santosh Ghosh, Avinash Varna, Majid Sabbagh, Sayak Ray, Rana Elnaggar, Joseph Friel, Daniel Dinu, Jason M. Fung</i>	
Discerning Limitations of GNN-Based Attacks on Logic Locking	1508
<i>Armin Darjani, Nima Kavand, Shubham Rai, Akash Kumar</i>	
Return-To-Non-Secure Vulnerabilities on ARM Cortex-M TrustZone: Attack and Defense	1514
<i>Zheyuan Ma, Xi Tan, Lukasz Ziarek, Ning Zhang, Hongxin Hu, Ziming Zhao</i>	
Automatic End-To-End Joint Optimization for Kernel Compilation on DSPs	1520
<i>Xiaolei Zhao, Zhaoyun Chen, Yang Shi, Mei Wen, Chunyun Zhang</i>	
VITAL: Vision Transformer Neural Networks for Accurate Smartphone Heterogeneity Resilient Indoor Localization	1526
<i>Danish Gufran, Saideep Tiku, Sudeep Pasricha</i>	
RLAlloc: A Deep Reinforcement Learning-Assisted Resource Allocation Framework for Enhanced Both I/O Throughput and QoS Performance of Multi-Streamed SSDs	1532
<i>Mengquan Li, Chao Wu, Congming Gao, Cheng Ji, Kenli Li</i>	
Towards Efficient Convolutional Neural Network for Embedded Hardware Via Multi-Dimensional Pruning	1538
<i>Hao Kong, Di Liu, Xiangzhong Luo, Shuo Huai, Ravi Subramaniam, Christian Makaya, Qian Lin, Weichen Liu</i>	

HTVM: Efficient Neural Network Deployment on Heterogeneous TinyML Platforms.....	1544
<i>Josse Van Delm, Maarten Vandersteegen, Alessio Burrello, Giuseppe Maria Sarda, Francesco Conti, Daniele Jahier Pagliari, Luca Benini, Marian Verhelst</i>	
Privacy-Preserving DNN Training with Prefetched Meta-Keys on Heterogeneous Neural Network Accelerators.....	1550
<i>Qiushi Li, Ju Ren, Yan Zhang, Chengru Song, Yiqiao Liao, Yaoxue Zhang</i>	
Safe DNN-Type Controller Synthesis for Nonlinear Systems Via Meta Reinforcement Learning	1556
<i>Hanrui Zhao, Xia Zeng, Niuniu Qi, Zhengfeng Yang, Zhenbing Zeng</i>	
GAMORA: Graph Learning Based Symbolic Reasoning for Large-Scale Boolean Networks.....	1562
<i>Nan Wu, Yingjie Li, Cong Hao, Steve Dai, Cunxi Yu, Yuan Xie</i>	
GenFuzz: GPU-Accelerated Hardware Fuzzing Using Genetic Algorithm with Multiple Inputs	1568
<i>Dian-Lun Lin, Yanqing Zhang, Haoxing Ren, Brucek Khailany, Shih-Hsin Wang, Tsung-Wei Huang</i>	
Partition Based Differential Testing for Finding Embedded Code Generation Bugs in Simulink.....	1574
<i>He Jiang, Hongyi Cheng, Shikai Guo, Xiaochen Li</i>	
Mantra: Mutation Testing of Hardware Design Code Based on Real Bugs.....	1580
<i>Jiang Wu, Yan Lei, Zhuo Zhang, Xiankai Meng, Deheng Yang, Pan Li, Jiayu He, Xiaoguang Mao</i>	
General-Purpose Gate-Level Simulation with Partition-Agnostic Parallelism.....	1586
<i>Zizheng Guo, Zuodong Zhang, Xun Jiang, Wuxi Li, Yibo Lin, Runsheng Wang, Ru Huang</i>	
Lightning Talk: Trinity - Assured Neuro-symbolic Model Inspired by Hierarchical Predictive Coding	1592
<i>Susmit Jha</i>	
Memory-Efficient and Real-Time SPAD-based dToF Depth Sensor with Spatial and Statistical Correlation.....	1594
<i>Shiyao Li, Zhenhua Zhu, Yu Zhu, Qingpeng Zhu, Jiangwei Zhang, Wenxiu Sun, Guohao Dai, Fei Qiao, Huazhong Yang, Yu Wang</i>	
Work Or Sleep: Freshness-Aware Energy Scheduling for Wireless Powered Communication Networks with Interference Consideration.....	1600
<i>Ling Lin, Lei Ju, Chun Jason Xue, Mingliang Zhou, Wei Zhang, Zimeng Zhou</i>	
APP: Enabling Soft Real-Time Execution on Densely-Populated Hybrid Memory System.....	1606
<i>Zheng-Wei Wu, Yun-Chih Chen, Yuan-Hao Chang, Tei-Wei Kuo</i>	
Shoggoth: Towards Efficient Edge-Cloud Collaborative Real-Time Video Inference Via Adaptive Online Learning.....	1612
<i>Liang Wang, Kai Lu, Nan Zhang, Xiaoyang Qu, Jianzong Wang, Jiguang Wan, Guokuan Li, Jing Xiao</i>	
Contention-Free Configured Grant Scheduling for 5G URLLC Traffic.....	1618
<i>Tianyu Zhang, Xiaobo Sharon Hu, Song Han</i>	
SEO: Safety-Aware Energy Optimization Framework for Multi-Sensor Neural Controllers at the Edge.....	1624
<i>Mohanad Odema, James Ferlez, Yasser Shoukry, Mohammad Abdullah Al Faruque</i>	

INVITED: Generalizing the ISA to the ILA: A Software/Hardware Interface for Accelerator-Rich Platforms	1630
<i>Bo-Yuan Huang, Hongce Zhang, Aarti Gupta, Sharad Malik</i>	
Invited: A Scalable Formal Approach for Correctness-Assured Hardware Design	1634
<i>Jin Yang, Jeremy Casas, Zhenkun Yang</i>	
Invited: Algorithms and Architectures for Accelerating Long Read Sequence Analysis.....	1638
<i>Hasindu Gamaarachchi, Kisaru Liyanage, Sri Parameswaran</i>	
Invited: Accelerating Genome Analysis Via Algorithm-Architecture Co-Design	1642
<i>Onur Mutlu, Can Firtina</i>	
Invited: Pre-Silicon Side Channel and Fault Analysis.....	1650
<i>Jasper Van Woudenberg, Peter Grossmann, Avinash L. Varna, Joseph Friel, Daniel Dinu, Ronnie Lindsay, Steve J. Brown</i>	
Invited: Caching in Automated Data Centric Vehicles for Edge Computing Scenarios	1654
<i>Nora Sperling, Alex Bendrick, Dominik Stöhrmann, Rolf Ernst</i>	
Invited: Waving the Double-Edged Sword: Building Resilient CAVs with Edge and Cloud Computing.....	1658
<i>Xiangguo Liu, Yunpeng Luo, Anthony Goeckner, Trishna Chakraborty, Ruochen Jiao, Ningfei Wang, Yixuan Wang, Takami Sato, Qi Alfred Chen, Qi Zhu</i>	
Invited: Automated Code Generation for Information Technology Tasks in YAML Through Large Language Models	1662
<i>Saurabh Pujar, Luca Buratti, Xiaojie Guo, Nicolas Dupuis, Burn Lewis, Sahil Suneja, Atin Sood, Ganesh Nalawade, Matt Jones, Alessandro Morari, Ruchir Puri</i>	
Invited: Compute Express Link™ (CXL™): An Open Interconnect for Cloud Infrastructure	1666
<i>Debendra Das Sharma</i>	
Invited: Building Robust Quantum System Software for Technology-Specific Characteristics	1670
<i>Tirthak Patel, Devesh Tiwari</i>	
CUDA Quantum: The Platform for Integrated Quantum-Classical Computing.....	1674
<i>Jin-Sung Kim, Alex McCaskey, Bettina Heim, Manish Modani, Sam Stanwyck, Timothy Costa</i>	
(Invited) Predictive Analytics for Cryogenic CMOS in Future Quantum Computing Systems	1678
<i>Rajiv V. Joshi, Sudipto Chakraborty</i>	
Invited: Algorithm-Software-Hardware Co-Design for Deep Learning Acceleration	1682
<i>Zhengang Li, Yanyue Xie, Peiyan Dong, Olivia Chen, Yanzhi Wang</i>	
INVITED: Design Automation Needs for Monolithic 3D ICs: Accomplishments and Gaps.....	1686
<i>Lingjun Zhu, Sung Kyu Lim</i>	
Invited: Autonomous Driving Digital Twin Empowered Design Automation: An Industry Perspective.....	1690
<i>Bo Yu, Jie Tang, Shao-Shan Liu</i>	
Invited: Buried Power Rails and Back-Side Power Grids: Prospects and Challenges.....	1694
<i>S. S. Teja Nibhanupudi, Sirish Oruganti, Rahul Mathur, Nishant Gupta, Meizhi Wang, Jaydeep P. Kulkarni</i>	

Late Breaking Results: Scalable and Efficient Hyperdimensional Computing for Network Intrusion Detection	1698
<i>Junyao Wang, Hanning Chen, Mariam Issa, Sitao Huang, Mohsen Imani</i>	
Late Breaking Results: Configurable Ring Oscillators as a Side-Channel Countermeasure	1700
<i>Hassan Nassar, Simon Pankner, Lars Bauer, Jörg Henkel</i>	
Late Breaking Results: PVT-Sensitive Delay Fitting for High-Performance Computing	1706
<i>Ding-Hao Wang, Shuo-Hung Hsu, Shu-Hsiang Yang, Pei-Ju Lin, Hui-Ting Yang, Mark Po-Hung Lin</i>	
Late Breaking Results: COPPER: Computation Obfuscation by Producing Permutations for Encoding Randomly	1708
<i>Kevin Hutto, Vincent Mooney</i>	
Late Breaking Results: Test Selection for RTL Coverage by Unsupervised Learning from Fast Functional Simulation	1710
<i>Rongjian Liang, Nathaniel Pinckney, Yuji Chai, Haoxin Ren, Brucek Khailany</i>	
Late Breaking Results: Weight Decay is ALL You Need for Neural Network Sparsification	1712
<i>Xizi Chen, Rui Pan, Xiaomeng Wang, Fengshi Tian, Chi-Ying Tsui</i>	
Late Breaking Results: RL-LPO: Reinforcement Learning Based Leakage Power Optimization Framework with Graph Neural Network	1714
<i>Peng Cao, Jiahao Wang</i>	
Late Breaking Results: An Efficient Bridge-Based Compression Algorithm for Topologically Quantum Error Corrected Circuits	1716
<i>Wei-Hsiang Tseng, Yao-Wen Chang</i>	
Late Breaking Results: Analytical Placement for 3D ICs with Multiple Manufacturing Technologies	1718
<i>Yan-Jen Chen, Yan-Syuan Chen, Wei-Che Tseng, Cheng-Yu Chiang, Yu-Hsiang Lo, Yao-Wen Chang</i>	
PIE-DRAM: Postponing IECC to Enhance DRAM Performance with Access Table	1720
<i>Jaehwa Jeon, Jae-Youn Hong, Sunghoon Kim, Insu Choi, Joon-Sung Yang</i>	
RQ-DNN: Reliable Quantization for Fault-tolerant Deep Neural Networks	1722
<i>Insu Choi, Jae-Youn Hong, Jaehwa Jeon, Joon-Sung Yang</i>	
Late Breaking Results from Hybrid Design Automation for Field-Coupled Nanotechnologies	1724
<i>Simon Hofmann, Marcel Walter, Lorenzo Servadei, Robert Wille</i>	
PVC-RAM: Process Variation Aware Charge Domain In-Memory Computing 6T-SRAM for DNNs	1726
<i>Sai Shubham, Shubham Pandit, Kailash Prasad, Joycee Mekie</i>	
Late Breaking Results: PyAIE: A Python-Based Programming Framework for Versal ACAP Platforms	1728
Late Breaking Results: Fast Fair Medical Applications? Hybrid Vision Models Achieve the Fairness on the Edge	1730
<i>Changdi Yang, Yi Sheng, Peiyan Dong, Zhenglun Kong, Yanyu Li, Pinrui Yu, Lei Yang, Xue Lin</i>	

Author Index