

2025 62nd ACM/IEEE Design Automation Conference (DAC 2025)

**San Francisco, California, USA
22-25 June 2025**

Pages 1-730



**IEEE Catalog Number: CFP25DAC-POD
ISBN: 979-8-3315-0305-5**

**Copyright © 2025 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:	CFP25DAC-POD
ISBN (Print-On-Demand):	979-8-3315-0305-5
ISBN (Online):	979-8-3315-0304-8

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

IntraFuzz: Coverage-Guided Intra-Enclave Fuzzing for Intel SGX Applications	1
<i>Jinhua Cui, Qiao Peng, Yiwen Yao, Ke Ye, Jiliang Zhang</i>	
BPUFuzzer: Effective Fuzz Testing for Branching Transient Execution Vulnerabilities of RISC-V CPU	8
<i>Rihui Sun, Jin Wu, Hanyin Liu, Zikang Tao, Gang Qu, Dongsheng Wang, Yongqiang Lyu, Jian Dong</i>	
ADVeRL-ELF: ADVersarial ELF Malware Generation Using Reinforcement Learning.....	15
<i>Akshara Ravi, Vivek Chaturvedi, Muhammad Shafique</i>	
Identifying System-On-Chip Security Assets with Structure-Based Analysis.....	22
<i>Wei-Kai Liu, Benjamin Tan, Krishnendu Chakraborty</i>	
ZION: A Practical Confidential Virtual Machine Architecture on Commodity RISC-V Processors.....	29
<i>Jie Wang, Juan Wang, Yinqian Zhang</i>	
FPGA-TrustZone: Security Extension of TrustZone to FPGA for SoC-FPGA Heterogeneous Architecture	36
<i>Shupeng Wang, Xindong Fan, Xiao Xu, Shuchen Wang, Lei Ju, Zimeng Zhou</i>	
MambaOPU: An FPGA Overlay Processor for State-Space-Duality-Based Mamba Models	42
<i>Shaoqiang Lu, Xuliang Yu, Tiandong Zhao, Siyuan Miao, Xinsong Sheng, Chen Wu, Liang Zhao, Ting-Jung Lin, Lei He</i>	
A Cross-Model Fusion-Aware Framework for Optimizing (gather-Matmul-Scatter)s Workload	49
<i>Yaoxiu Lian, Zhihong Gou, Yibo Han, Zhongming Yu, Jiaming Xu, Sheng Yuan, Zhilin Pei, Xingcheng Zhang, Ningyi Xu, Guohao Dai</i>	
HybriMoE: Hybrid CPU-GPU Scheduling and Cache Management for Efficient MoE Inference.....	56
<i>Shuzhang Zhong, Yanfan Sun, Ling Liang, Runsheng Wang, Ru Huang, Meng Li</i>	
ClusterKV: Manipulating LLM KV Cache in Semantic Space for Recallable Compression.....	63
<i>Guangda Liu, Chengwei Li, Jieru Zhao, Chenqi Zhang, Minyi Guo</i>	
DuoQ: A DSP Utilization-Aware and Outlier-Free Quantization for FPGA-Based LLMs Acceleration.....	70
<i>Zhuoquan Yu, Huidong Ji, Yue Cao, Junfu Wu, Xiaoze Yan, Lirong Zheng, Zhuo Zou</i>	
Libra: A Hybrid-Sparse Attention Accelerator Featuring Multi-Level Workload Balance	77
<i>Faxian Sun, Runzhou Zhang, Zhenyu Liu, Heng Liao, Zhinan Qin, Jianli Chen, Jun Yu, Kun Wang</i>	
UniCAIM: A Unified CAM/CIM Architecture with Static-Dynamic KV Cache Pruning for Efficient Long-Context LLM Inference	84
<i>Weikai Xu, Wenxuan Zeng, Qianqian Huang, Meng Li, Ru Huang</i>	
P-DAC: Power-Efficient Photonic Accelerators for LLM Inference.....	91
<i>Wen-Tse Chang, Chun-Feng Wu, Yun-Chen Lo</i>	

A PulseWidth-IN-PulseWidth-Out Universal Nonlinear Processing Element for Time-Domain In-Memory Computing Systems	98
<i>Yihao Chen, Pengcheng Feng, Zhigang Li, Gang Chen, Rongxuan Shen, Huaxiang Lu, Xiaoxin Xu</i>	
PUFiM: A Robust and Efficient FeFET-Based Security Solution Merging Physical Unclonable Function with Compute-In-Memory for Edge AI.....	104
<i>Taixin Li, Thomas Kämpfe, Jianfeng Wang, Kai Ni, Vijaykrishnan Narayanan, Huazhong Yang, Xueqing Li</i>	
TAXI: Traveling Salesman Problem Accelerator with X-Bar-Based Ising Macros Powered by SOT-MRAMs and Hierarchical Clustering.....	111
<i>Sangmin Yoo, Amod Holla, Sourav Sanyal, Dong Eun Kim, Francesca Iacopi, Dwaipayan Biswas, James Myers, Kaushik Roy</i>	
A Mixed-Signal Photonic SRAM-Based High-Speed Energy-Efficient Photonic Tensor Core with Novel Electro-Optic ADC	118
<i>Md Abdullah-Al Kaiser, Sugeet Sunder, Ajey P. Jacob, Akhilesh R. Jaiswal</i>	
INSTA: An Ultra-Fast, Differentiable, Statistical Static Timing Analysis Engine for Industrial Physical Design Applications	125
<i>Yi-Chen Lu, Zhizheng Guo, Kishor Kunal, Rongjian Liang, Haoxing Ren</i>	
GTN-Path: Efficient Path Timing Prediction Through Waveform Propagation with Graph Transformer	132
<i>Lihao Liu, Beisi Lu, Yunhui Li, Li Shang, Fan Yang</i>	
Generative Model Based Standard Cell Timing Library Characterization	139
<i>Hao-Yu Wu, Hsin-Tzu Chang, Shiuuan-Yun Ding, Iris Hui-Ru Jiang, Benson Tsao, Vinson Wu, Wei-Kai Shih</i>	
Truly Pre-Routing Timing Prediction Via Considering Power Delivery Network	146
<i>Yuyang Ye, Mingwei He, Lizheng Ren, Jianwang Zhai, Tinghuan Chen, Jun Yang, Longxing Shi</i>	
A Fast, Iterative Clock Skew Scheduling Algorithm with Dynamic Sequential Graph Extraction	153
<i>Shijian Chen, Yihang Qiu, Biwei Xie, Mingyu Chen, Xingquan Li</i>	
GLiTCH: GLiTCH Induced Transitions for Secure Crypto-Hardware	160
<i>C. Rohin Menon, Jayanth Balasubramanian, E Akshay Kumar, Annapurna Valiveti, Chester Rebeiro, Janakiraman Viraraghavan</i>	
AmpereBleed: Exploiting On-Chip Current Sensors for Circuit-Free Attacks on ARM-FPGA SoCs	167
<i>Xin Zhang, Yi Yang, Jiajun Zou, Qingni Shen, Zhi Zhang, Yansong Gao, Zhonghai Wu, Trevor E. Carlson</i>	
ZenLeak: Practical Last-Level Cache Side-Channel Attacks on AMD Zen Processors	174
<i>Han Wang, Ming Tang, Quancheng Wang, Ke Xu, Yinqian Zhang</i>	
ZK-Hammer: Leaking Secrets from Zero-Knowledge Proofs Via Rowhammer.....	181
<i>Junkai Liang, Xin Zhang, Daqi Hu, Qingni Shen, Yuejian Fang, Zhonghai Wu</i>	
Cross-Attention for AES Mode Variation in Side-Channel Analysis.....	188
<i>Fanliang Hu, Jian Shen, Haoyu Ma, Qingming Jonathan Wu</i>	
Security of Approximate Neural Networks Against Power Side-Channel Attacks.....	195
<i>Aditya Japa, Jack Miskelly, Maire O'Neill, Chongyan Gu</i>	

TetrisLock: Quantum Circuit Split Compilation with Interlocking Patterns	202
<i>Qian Wang, Jayden John, Ben Dong, Yuntao Liu</i>	
Towards Training Robustness Against Dynamic Errors in Quantum Machine Learning	209
<i>Shijin Duan, Gaowen Liu, Charles Fleming, Ramana Kompella, Xiaolin Xu, Shaolei Ren</i>	
CirSTAG: Circuit Stability Analysis on Graph-Based Manifolds	216
<i>Wuxinlin Cheng, Yihang Yuan, Chenhui Deng, Ali Aghdaei, Zhiru Zhang, Zhuo Feng</i>	
ParGNN: A Scalable Graph Neural Network Training Framework on multi-GPUs.....	223
<i>Junyu Gu, Shunde Li, Rongqiang Cao, Jue Wang, Zijian Wang, Zhiqiang Liang, Fang Liu, Shigang Li, Chunbao Zhou, Yangang Wang, Xuebin Chi</i>	
Delving into Topology Representation for Layout Pattern: A Novel Contrastive Learning Framework for Hotspot Detection.....	230
<i>Silin Chen, Kangjian Di, Guohao Wang, Wenzheng Zhao, Li Du, Ningmu Zou</i>	
SuperFast: Fast Supernet Training Using Initial Knowledge	236
<i>Moritz Thoma, Emad Aghajanzadeh, Shambhavi Balamuthu Sampath, Pierpaolo Mori, Nael Fafous, Alexander Frickenstein, Manoj-Rohit Vemparala, Daniel Mueller-Gritschneider, Ulf Schlichtmann</i>	
LA-MTL: Latency-Aware Automated Multi-Task Learning	243
<i>Shambhavi Balamuthu Sampath, Sami Sawani, Moritz Thoma, Lukas Frickenstein, Pierpaolo Mori, Nael Fafous, Manoj Rohit Vemparala, Alexander Frickenstein, Ulf Schlichtmann, Claudio Passerone, Walter Stechele</i>	
GLOVA: Global and Local Variation-Aware Analog Circuit Design with Risk-Sensitive Reinforcement Learning.....	250
<i>Dongjun Kim, Junwoo Park, Chaehyeon Shin, Jaeheon Jung, Kyungho Shin, Seungheon Baek, Sanghyuk Heo, Woongrae Kim, Inchul Jeong, Joohwan Cho, Jongsun Park</i>	
Graph-Guided Transfer Learning to Boost the Efficiency of System-Level Optimization of Analog/Mixed-Signal Circuits.....	257
<i>Jiaqi Wang, Georges G. E. Gielen</i>	
INSIGHT: A Universal Neural Simulator Framework for Analog Circuits with Autoregressive Transformers.....	264
<i>Souradip Poddar, Youngmin Oh, Yao Lai, Hanqing Zhu, Bosun Hwang, David Z. Pan</i>	
G-SpNN: GPU-Accelerated Passivity Enforcement for S-Parameter Modeling with Neural Networks	271
<i>Lijie Zeng, Jiatai Sun, Xiao Wu, Dan Niu, Tianshi Wang, Yibo Lin, Zuochang Ye, Zhou Jin</i>	
A Novel Image-Graph Heterogeneous Fusion Framework for Static IR Drop Prediction	278
<i>Dan Niu, Dekang Zhang, Yichao Cao, Zhou Jin, Chao Wang, Yichao Dong, Changyin Sun</i>	
EPOC: An Efficient Pulse Generation Framework with Advanced Synthesis for Quantum Circuits.....	285
<i>Jinglei Cheng, Yuchen Zhu, Yidong Zhou, Hang Ren, Zhixin Song, Zhiding Liang</i>	
Assessing Quantum Layout Synthesis Tools Via Known Optimal-SWAP Cost Benchmarks	292
<i>Shuohao Ping, Wan-Hsuan Lin, Daniel Bochen Tan, Jason Cong</i>	
A Scalable and Robust Compilation Framework for Emitter-Photonic Graph State	299
<i>Xiangyu Ren, Yuexun Huang, Zhiding Liang, Antonio Barbalace</i>	

PHOENIX: Pauli-Based High-Level Optimization Engine for Instruction Execution on NISQ Devices	306
<i>Zhaohui Yang, Dawei Ding, Chenghong Zhu, Jianxin Chen, Yuan Xie</i>	
Joint Cutting for Hybrid Schrödinger-Feynman Simulation of Quantum Circuits	313
<i>Laura S. Herzog, Lukas Burgholzer, Christian Ufrecht, Daniel D. Scherer, Robert Wille</i>	
Versatile Cross-Platform Compilation Toolchain for Schrödinger-Style Quantum Circuit Simulation.....	320
<i>Yuncheng Lu, Shuang Liang, Hongxiang Fan, Ce Guo, Wayne Luk, Paul H. J. Kelly</i>	
LVM-MO: A Large Vision Model Pioneer on Full-Chip Mask Optimization	327
<i>Yiwen Wu, Yuyang Chen, Shuo Yin, Nan Wang, Tao Wu, Xuming He, Hao Geng, Jingyi Yu</i>	
SSDL-ILT: Efficient ILT Utilizing a Self-Supervised Deep Learning Model	334
<i>Rui Xu, Junqi Yang, Haoxiang Jiang, Ming Fang</i>	
PatternPaint: Practical Layout Pattern Generation Using Diffusion-Based Inpainting	340
<i>Guanglei Zhou, Bhargav Korrapati, Gaurav Rajavendra Reddy, Chen-Chia Chang, Jingyu Pan, Jiang Hu, Yiran Chen, Dipto G. Thakurta</i>	
Generalizable Lithographic Hotspot Detection Using Asynchronous Meta-Learning with Only One Shot.....	347
<i>Cong Jiang, Yujia Wang, Dan Feng, Haoyu Yang, Kang Liu</i>	
Accuracy is Not Always We Need: Precision-Aware Bayesian Yield Optimization	354
<i>Jing Kou, Zidong Chen, Liang Zhang, Haiyan Qin, Wang Kang, Wei W. Xing</i>	
Curvilinear Optical Proximity Correction Via Cardinal Spline	361
<i>Su Zheng, Xiaoxiao Liang, Ziyang Yu, Yuzhe Ma, Bei Yu, Martin Wong</i>	
MIRACLE: Multimodal Information Retrieval Via a Combined In-Memory Processing and Content Addressable Memory Approach.....	368
<i>Xuehui Liu, Xueyan Wang, Tianyang Yu, Chen Cheng, Shuo Ran, Bi Wu, Xiaotao Jia, Weiqiang Liu, Gang Qu, Weisheng Zhao</i>	
REMU: Memory-Aware Radiation Emulation Via Dual Addressing for In-Orbit Deep Learning System.....	375
<i>Longnv Xu, Meiqi Wang, Han Qiu, Jun Liu, Yuanjie Li, Hewu Li</i>	
HIVE: A High-Priority Victim Cache for Accelerating GPU Memory Accesses	382
<i>Yuhan Tang, Jianmin Zhang, Sheng Ma, Tiejun Li, Hanqing Li, Shengbai Luo, Jixuan Tang, Lizhou Wu</i>	
NVR: Vector Runahead on NPUs for Sparse Memory Access.....	389
<i>Hui Wang, Zhengpeng Zhao, Jing Wang, Yushu Du, Yuan Cheng, Bing Guo, He Xiao, Chenhao Ma, Xiaomeng Han, Dean You, Jiapeng Guan, Ran Wei, Dawei Yang, Zhe Jiang</i>	
Expanding Logical Space Freely: A Memory-Efficient Mapping Table Design for Compressional SSDs.....	396
<i>Zixuan Huang, Tianyu Wang, Kecheng Huang, Zelin Du, Zili Shao</i>	
Sphinx: A High-Performance Hybrid Index for Disaggregated Memory with Succinct Filter Cache.....	403
<i>Jingxiang Li, Shengan Zheng, Bowen Zhang, Hankun Dong, Linpeng Huang</i>	
Ragnar: Exploring Volatile-Channel Vulnerabilities on RDMA NIC	410
<i>Yunpeng Xu, Yuchen Fan, Teng Ma, Shuwen Deng</i>	

Data Oblivious CPU: Microarchitectural Side-Channel Leakage-Resilient Processor	417
<i>Behnam Omidi, Ihsen Alouani, Khaled N. Khasawneh</i>	
“OOPSI!”: Out-Of-Band Remote Power Side-Channel Attacks on Intel SGX and TDX.....	424
<i>Nimish Mishra, Kislay Arya, Sarani Bhattacharya, Paritosh Saxena, Debdeep Mukhopadhyay</i>	
POLARIS: Explainable Artificial Intelligence for Mitigating Power Side-Channel Leakage.....	431
<i>Tanzim Mahfuz, Sudipta Paria, Tasneem Suha, Swarup Bhunia, Prabuddha Chakraborty</i>	
SCONE: A Logic Locking Technique Utilizing SMT Solver and Circuit Encoding Scheme for Efficient Hardware IP Protection.....	438
<i>Zhaokun Han, Daniel Xing, Kostas Amberiadis, Ankur Srivastava, Jeyavijayan Jv Rajendran</i>	
DeepPUFSCA: Deep Learning for Physical Unclonable Function Attack Based on Side Channel Analysis Support	445
<i>Ngoc Phu Doan, Tuan Dung Pham, Zichi Zhang, Viet Hung Tran, Jack Miskelly, Hans Vandierendonck, Anh Tuan Hoang, Máire O’Neill, Thai Son Mai</i>	
Grasp: Group-Based Prediction of Activation Sparsity for Fast LLM Inference	452
<i>Jiho Shin, Hoeseok Yang, Youngmin Yi</i>	
DuQTTA: Dual Quantized Tensor-Train Adaptation with Decoupling Magnitude-Direction for Efficient Fine-Tuning of LLMs	459
<i>Haoyan Dong, Hai-Bao Chen, Jingjing Chang, Yixin Yang, Ziyang Gao, Zhigang Ji, Runsheng Wang, Ru Huang</i>	
PacTrain: Pruning and Adaptive Sparse Gradient Compression for Efficient Collective Communication in Distributed Deep Learning.....	466
<i>Yisu Wang, Ruilong Wu, Xinjiao Li, Dirk Kutscher</i>	
MILLION: MasterIng Long-Context LLM Inference Via Outlier-Immunized KV Product QuaNtization	473
<i>Zongwu Wang, Peng Xu, Fangxin Liu, Yiwei Hu, Qingxiao Sun, Gezi Li, Cheng Li, Xuan Wang, Li Jiang, Haibing Guan</i>	
AASD: Accelerate Inference by Aligning Speculative Decoding in Multimodal Large Language Models.....	480
<i>Chaoqun Yang, Ran Chen, Muyang Zhang, Weiguang Pang, Yuzhi Chen, Rongtao Xu, Kexue Fu, Changwei Wang, Longxiang Gao</i>	
ChipAlign: Instruction Alignment in Large Language Models for Chip Design Via Geodesic Interpolation	487
<i>Chenhui Deng, Yunsheng Bai, Haoxing Ren</i>	
FeKAN: Efficient Kolmogorov-Arnold Networks Accelerator Using FeFET-Based CAM and LUT	494
<i>Xuliang Yu, Yu Qian, Xunzhao Yin, Cheng Zhuo, Liang Zhao</i>	
Energy-Efficient Large-Scale Vector Similarity Search in NAND-Flash Via Hybrid Matching.....	501
<i>Chih-Yu Hu, Chi-Tse Huang, Hao-Wei Chiang, Hsiang-Yun Cheng, Po-Hao Tseng, Ming-Hsiu Lee, An-Yeu Andy Wu</i>	
Chameleon-SAT: An Adaptive Boolean Satisfiability Accelerator Using Mixed-Signal In-Memory Computing for Versatile SAT Problems.....	508
<i>Iris Ying Chou, Hao Kong, Yi Huang, Jianfeng Zhu, Wenping Zhu, Shaojun Wei, Aoyang Zhang, Leibo Liu</i>	

A Full-System, Programmable, and Extensible In-Memory Computing Simulation Framework for Deep Learning	515
<i>Kaining Zhou, Jian Huang, Nam Sung Kim, Naresh Shanbhag</i>	
All-In-Memory Stochastic Computing Using ReRAM	522
<i>João Paulo C. De Lima, Mehran Shoushtari Moghadam, Sercan Aygun, Jeronimo Castrillon, M. Hassan Najafi, Asif Ali Khan</i>	
UPVSS: Jointly Managing Vector Similarity Search with Near-Memory Processing Systems.....	529
<i>Chun-Chien Liu, Chun-Feng Wu, Yunho Jin</i>	
SSFT: Algorithm and Hardware Co-Design for Structured Sparse Fine-Tuning of Large Language Models	536
<i>Miao Yu, Trevor E. Carlson</i>	
DRAFT: Decoupling Backpropagation from Pre-Trained Backbone for Efficient Transformer Fine-Tuning on Edge	543
<i>Zhirui Huang, Shiwei Liu, Haozhe Zhu, Qi Liu, Chixiao Chen</i>	
Move Less, Retrieve Fast: A Retrieval-In-Memory Architecture for Language Models.....	550
<i>Jiaxian Chen, Yuxuan Qi, Jianan Yuan, Kaoyi Sun, Tianyu Wang, Chenlin Ma, Yi Wang</i>	
KVO-LLM: Boosting Long-Context Generation Throughput for Batched LLM Inference.....	557
<i>Zhenyu Li, Dongxu Lyu, Gang Wang, Yuzhou Chen, Liyan Chen, Wenjie Li, Jianfei Jiang, Yanan Sun, Guanghui He</i>	
An Energy-Efficient High-Utilization Hardware Architecture for Attention Mechanism in Transformer Using Balanced Systolic Array and Multi-Row Interleaved Operation Ordering.....	564
<i>Haiyang Zhou, Hongyang Hu, Jinshan Yue, Hanghang Gao, Yuanlu Xie, Xiaoxin Xu, Chunmeng Dou, Ming Liu</i>	
Finding the Pareto Frontier of Low-Precision Data Formats and MAC Architecture for LLM Inference	571
<i>Brian Crafton, Xiaochen Peng, Xiaoyu Sun, Ashwin Lele, Bo Zhang, Win-San Khwa, Kerem Akarvardar</i>	
Fast End-To-End Simulation and Exploration of Many-RISCV-Core Baseband Transceivers for Software-Defined Radio-Access Networks	578
<i>Marco Bertuletti, Yichao Zhang, Mahdi Abdollahpour, Samuel Riedel, Alessandro Vanelli-Coralli, Luca Benini</i>	
Intermittent Systems at Small Scale: Execution Model and Design Guidelines	585
<i>Youngbin Kim, Yoojin Lim</i>	
MEEK: Re-Thinking Heterogeneous Parallel Error Detection Architecture for Real-World OoO Superscalar Processors	592
<i>Zhe Jiang, Minli Liao, Sam Ainsworth, Dean You, Timothy Jones</i>	
VersaSlot: Efficient Fine-Grained FPGA Sharing with Big.Little Slots and Live Migration in FPGA Cluster	599
<i>Jianfeng Gu, Hao Wang, Xiaorang Guo, Martin Schulz, Michael Gerndt</i>	
Power-Constrained Printed Neuromorphic Hardware Training.....	606
<i>Tara Gheshlaghi, Haibin Zhao, Priyanjana Pal, Michael Hefenbrock, Michael Beigl, Mehdi B. Tahoori</i>	

Age-Of-Information Minimization for Data Aggregation in Energy-Harvesting IoTs	613
<i>Bingkun Yao, Mun Choon Chan, Hong Gao, Zhe Jiang, Nan Guan</i>	
Uncertainty-Aware Energy Management for Wearable IoT Devices with Conformal Prediction	620
<i>Dina Hussein, Chibuikwe E. Ugwu, Ganapati Bhat, Janardhan Rao Doppa</i>	
Query-Based Black-Box Stealthy Sensor Attacks on Cyber-Physical Systems	627
<i>Shixiong Jiang, Weizhe Xu, Mengyu Liu, Fanxin Kong</i>	
Pipirima: Predicting Patterns in Sparsity to Accelerate Matrix Algebra.....	633
<i>Ubaid Bakhtiar, Donghyeon Joo, Bahar Asgari</i>	
FactorHD: A Hyperdimensional Computing Model for Multi-Object Multi-Class Representation and Factorization	640
<i>Yifei Zhou, Xuchu Huang, Chenyu Ni, Min Zhou, Zheyu Yan, Xunzhao Yin, Cheng Zhuo</i>	
Holistic Design Towards Resource-Stringent Binary Vector Symbolic Architecture	647
<i>Shijin Duan, Nuntipat Narkthong, Yukui Luo, Shaolei Ren, Xiaolin Xu</i>	
SDISC: A Spike-Driven Human-Machine Interface with In-Situ Computing for Real-Time Low- Power Interaction	654
<i>Fangduo Zhu, Jingyi Chen, Jingsong Zhang, Xumeng Zhang, Siyuan Ouyang, Chenyang, Hao Jiang, Xiaonan Yang, Qi Liu</i>	
PairGraph: An Efficient Search-Space-Aware Accelerator for High-Performance Concurrent Pairwise Queries	661
<i>Yutao Fu, Zhongtian Long, Yu Zhang, Zirui He, Jin Zhao, Qiyuan Niu, Zixiao Wang, Hai Jin</i>	
PreDAC: An Efficient Framework of Pre-Refining Enhanced Design Space Exploration for Approximate Computing	668
<i>Ziying Cui, Ke Chen, Bi Wu, Yu Gong, Chenggang Yan, Weiqiang Liu</i>	
AutoRE: Bayesian-Optimization-Based Automatic Reliability Enhancement Tool for Flow-Based Microfluidic Biochips.....	675
<i>Siyuan Liang, Yushen Zhang, Mengchu Li, Tsun-Ming Tseng, Ulf Schlichtmann, Tsung-Yi Ho</i>	
DDRRoute: A Novel Depth-Driven Approach to the Qubit Routing Problem.....	682
<i>Alessandro Annechini, Marco Venere, Donatella Sciuto, Marco D. Santambrogio</i>	
KLiNQ: Knowledge Distillation-Assisted Lightweight Neural Network for Qubit Readout on FPGA.....	689
<i>Xiaorang Guo, Tigran Bunarjyan, Dai Liu, Benjamin Lienhard, Martin Schulz</i>	
Efficient and Scalable Architectures for Multi-Level Superconducting Qubit Readout.....	696
<i>Chaithanya Naik Mude, Satvik Maurya, Benjamin Lienhard, Swamit Tannu</i>	
DyREM: Dynamically Mitigating Quantum Readout Error with Embedded Accelerator	703
<i>Kaiwen Zhou, Liqiang Lu, Hanyu Zhang, Debin Xiang, Chenning Tao, Xinkui Zhao, Size Zheng, Jianwei Yin</i>	
Weighted Range-Constrained Ising-Model Decoder for Quantum Error Correction	710
<i>Xinyi Guo, Hiromitsu Awano, Takashi Sato</i>	
ZXNet: ZX Calculus-Driven Graph Neural Network Framework for Quantum Circuit Equivalence Checking.....	717
<i>Navnil Choudhury, Ameya S Bhave, Kanad Basu</i>	

EnQode: Fast Amplitude Embedding for Quantum Machine Learning Using Classical Data	724
<i>Jason Han, Nicholas S. Dibrita, Younghyun Cho, Hengrui Luo, Tirthak Patel</i>	
ArbiterQ: Improving QNN Convergence and Accuracy by Applying Personalized Model on Heterogeneous Quantum Devices	731
<i>Tianyao Chu, Siwei Tan, Liqiang Lu, Jingwen Leng, Fangxin Liu, Congliang Lang, Yifan Guo, Jianwei Yin</i>	
Can Short Hypervectors Drive Feature-Rich GNNs? Strengthening the Graph Representation of Hyperdimensional Computing for Memory-Efficient GNNs	738
<i>Jihe Wang, Yuxi Han, Danghui Wang</i>	
InfScaler: Enabling Efficient ML Inference Serving on Multi-Accelerator Edge Devices Via Asymmetric Auto-Scaling	745
<i>Borui Li, Tiange Xia, Shuai Wang, Shuai Wang</i>	
DM-Tune: Quantizing Diffusion Models with Mixture-Of-Gaussian Guided Noise Tuning	752
<i>Pouya Haghi, Ali Falahati, Zahra Azad, Chunshu Wu, Ruibing Song, Chuan Liu, Ang Li, Tong Geng</i>	
Towards In-Situ Neuromorphic Computing Architecture for Event Stream Super-Resolution	759
<i>Yihe Yu, Bo Li, Kexin Huang, Wei Liu, Jinghai Wang, Yue Liu, Zhiyi Yu, Shanlin Xiao</i>	
LearnGraph: A Learning-Based Architecture for Dynamic Graph Processing	766
<i>Lingling Zhang, Yijian Wu, Hong Jiang, Ziyu Zhou, Tiancheng Lu</i>	
Scaling Laws of Graph Neural Networks for Atomistic Materials Modeling*	773
<i>Chaojian Li, Zhifan Ye, Massimiliano Lupo Pasini, Jong Youl Choi, Cheng Wan, Prasanna Balaprakash</i>	
VISTA: Optimizing GPU Scheduling Through Versatile Locality-Aware Data Sharing	780
<i>Hajar Falahati, Negin Mahani, Adrian Cristal, Osman Unsal</i>	
Tropical: Enhancing SLO Attainment in Disaggregated LLM Serving Via SLO-Aware Multiplexing	787
<i>Jinming Ma, Jiefei Chen, Xiuhong Li, Jiangfei Duan, Haojie Duanmu, Xingcheng Zhang, Chao Yang, Dahua Lin</i>	
DBC: Drift-Aware Binary Code for Drift-Tolerant Deep Neural Networks	794
<i>Insu Choi, Jaeyong Chung, Joon-Sung Yang</i>	
CXL-INTERPLAY: Unraveling and Characterizing CXL Interference in Modern Computer Systems	801
<i>Shunyu Mao, Jiajun Luo, Yixin Li, Jiapeng Zhou, Weidong Zhang, Zheng Liu, Teng Ma, Shuwen Deng</i>	
VQT-CiM: Accelerating Vector Quantization Enhanced Transformer with Ferroelectric Compute-In-Memory	808
<i>Xuchu Huang, Haonan Du, Min Zhou, Zheyu Yan, Cheng Zhuo, Xunzhao Yin</i>	
Device-Algorithm Co-Design of Ferroelectric Compute-In-Memory In-Situ Annealer for Combinatorial Optimization Problems	815
<i>Yu Qian, Xianmin Huang, Ranran Wang, Zeyu Yang, Min Zhou, Thomas Kämpfe, Cheng Zhuo, Xunzhao Yin</i>	
DANN: Diffractive Acoustic Neural Network for In-Sensor Computing System Target at Multi-Biomarker Diagnosis	822
<i>Lewei He, Ning Lin, Binbin Cui, Xinran Zhang, Shiming Zhang, Zhongrui Wang</i>	

333-EDRAM - 3T Embedded DRAM Leveraging Monolithic 3D Integration of 3 Transistor Types: IGZO, Carbon Nanotube and Silicon FETs	829
<i>David Kong, Shvetank Prakash, Jdrzej Kufel, Georgios Kyriazidis, Yasmine Omri, David Verity, Emre Ozer, Vijay Janapa Reddi, Gage Hills</i>	
Monolithic 3D FPGA Design and Synthesis with Back-End-Of-Line Configuration Memories	836
<i>Faaq Waqar, Jiahao Zhang, Anni Lu, Zifan He, Jason Cong, Shimeng Yu</i>	
CaMDN: Enhancing Cache Efficiency for Multi-Tenant DNNs on Integrated NPUs.....	843
<i>Tianhao Cai, Liang Wang, Limin Xiao, Meng Han, Zeyu Wang, Lin Sun, Xiaojian Liao</i>	
Configurable DSP-Based CAM Architecture for Data-Intensive Applications on FPGAs	850
<i>Yao Chen, Feng Yu, Di Wu, Weng-Fai Wong, Bingsheng He</i>	
HPIM-NoC: A Priori-Knowledge-Based Optimization Framework for Heterogeneous PIM-Based NoCs.....	857
<i>Shuai Yuan, Angxin Cai, Qiushi Lin, Guoxing Wang, Yu Wang, Zhenhua Zhu, Yanan Sun</i>	
McPAL: Scaling Unstructured Sparse Inference with Multi-Chiplet HBM-PIM Architecture for LLMs.....	864
<i>Shiwei Liu, Zhirui Huang, Jiangnan Yu, Qi Liu, Chixiao Chen</i>	
Hypnos: Memory Efficient Homomorphic Processing Unit.....	871
<i>Haoxuan Wang, Yinghao Yang, Hang Lu, Xiaowei Li</i>	
GS-TG: 3D Gaussian Splatting Accelerator with Tile Grouping for Reducing Redundant Sorting While Preserving Rasterization Efficiency.....	878
<i>Joongho Jo, Jongsun Park</i>	
BEVSA: A Real-Time Bird's-Eye-View Semantic Segmentation Accelerator for Multi-Camera System	885
<i>Sangho Lee, Jueung Jung, Wuyoung Jang, Jihyeon Hwang, Kyuho Lee</i>	
An Energy-Efficient Kalman Filter Architecture with Tunable Accuracy for Brain-Computer Interfaces	892
<i>Guy Eichler, Joseph Zuckerman, Luca P. Carloni</i>	
A Cutting-Edge Parallel Solver for Scalable Power Grid Analysis Using Nested Domain Decomposition.....	899
<i>Jianfei Song, Xiaoyu Yang, Zhou Jin, Cheng Zhuo</i>	
IRGNN: A Graph-Based Framework Integrating Numerical Solution and Point Cloud for Static IR Drop Prediction	906
<i>Feng Guo, Yueyue Xi, Jianwang Zhai, Jingyu Jia, Jiawei Liu, Kang Zhao, Chuan Shi</i>	
LMM-IR: Large-Scale Netlist-Aware Multimodal Framework for Static IR-Drop Prediction	913
<i>Kai Ma, Zhen Wang, Hongquan He, Qi Xu, Tinghuan Chen, Hao Geng</i>	
Power-Grid Structure Exploration with Unified Sequence-Based Learning Framework.....	920
<i>Yi-Lin Chuang, Hao-Wei Chan, Chih-Yun Yen, Shih-An Hsieh, Ching-Feng Chen, Sheng-Te Lai</i>	
Real-Time Dynamic IR-Drop Prediction for IR ECO	927
<i>Yu-Che Lee, Yu-Hsuan Chen, Yu-Chen Cheng, Yong-Fong Chang, Jia-Wei Lin, Hsun-Wei Pao, Yung-Chih Chen, Yi-Ting Li, Wuqian Tang, Shih-Chieh Chang, Chun-Yao Wang</i>	

ATLAS: A Self-Supervised and Cross-Stage Netlist Power Model for Fine-Grained Time-Based Layout Power Analysis.....	934
<i>Wenkai Li, Yao Lu, Wenji Fang, Jing Wang, Qijun Zhang, Zhiyao Xie</i>	
NeuralMesh: Neural Network for FEM Mesh Generation in 2.5D/3D Chiplet Thermal Simulation.....	941
<i>Pengju Chen, Dan Niu, Dekang Zhang, Wenhao Wang, Depeng Xie, Zhou Jin, Wei W. Xing, Lei He</i>	
ASRR-PINN: Adaptive Sub-Regional Random Resampling-Based PINN for Thermal Analysis of 3D-ICs.....	948
<i>Zijian Zhou, Min Tang, Liang Chen</i>	
Clearance-Constrained PCB Global Placement with Heterogeneous Components.....	955
<i>Yan-Jen Chen, Wei-Kai Huang, Chung-Ting Tsai, Chiao-Yu Ou, Yao-Wen Chang</i>	
Constraint Graph-Based PCB Legalization Considering Dense, Heterogeneous, Irregular-Shaped, and Any-Oriented Components	962
<i>Chiao-Yu Ou, Yan-Jen Chen, Yao-Wen Chang</i>	
GPart: A GNN-Enabled Multilevel Graph Partitioner.....	969
<i>Magi Chen, Ting-Chi Wang</i>	
BlasPart: A Deterministic Parallel Partitioner for Balanced Large-Scale Hypergraph Partitioning	976
<i>Shengbo Tong, Chunyan Pei, Wenjian Yu</i>	
MIA-Aware FinFlex Cell Legalization with Power-Driven Cell Version Substitution.....	983
<i>Da-Wei Huang, Shao-Yun Fang</i>	
Secondary-Power-Cell-Aware Detailed Placement in Multiple Power Domain Designs.....	990
<i>Yu-Wei Chang, Shao-Yun Fang, Kai-Chuan Yang, Min-Ching Lin</i>	
High Energy-Efficiency and Low Latency In-Memory Computing Using Analog Accumulator and In-Memory ADC with Shared References.....	997
<i>Junyi Yang, Shuai Dong, Zhengnan Fu, Hongyang Shang, Arindam Basu</i>	
CREST-CiM: Cross-Coupling-Enhanced Differential STT-MRAM for Robust Computing-In-Memory in Binary Neural Networks	1004
<i>Imtiaz Ahmed, Akul Malhotra, Sumeet Kumar Gupta</i>	
YOCO: A Hybrid In-Memory Computing Architecture with 8-Bit Sub-PetaOps/W In-Situ Multiply Arithmetic for Large-Scale AI.....	1011
<i>Zihao Xuan, Yuxuan Yang, Wei Xuan, Zijia Su, Song Chen, Yi Kang</i>	
ReSMiPS: A ReRAM-Based Sparse Mixed-Precision Solver with Fast Matrix Reordering Algorithm	1018
<i>Yuyang Fu, Jiancong Li, Jia Chen, Zhiwei Zhou, Houji Zhou, Wenlong Peng, Yi Li, Xiangshui Miao</i>	
Lookup Table-Based Multiplication-Free All-Digital DNN Accelerator Featuring Self-Synchronous Pipeline Accumulation.....	1025
<i>Hiroto Tagata, Takashi Sato, Hiromitsu Awano</i>	
WISEDRAM: A Reliable Bitwise In-DRAM Accelerator.....	1032
<i>Mohammad Arman Soleimani, Nezam Rohbani, Adrian Cristal Kestelman, Osman Unsal, Hamid Sarbazi-Azad</i>	

Resilient Federated Learning on Embedded Devices with Constrained Network Connectivity	1039
<i>Zihan Li, Han Liu, Ao Li, Ching-Hsiang Chan, Yevgeniy Vorobeychik, William Yeoh, Wenjing Lou, Ning Zhang</i>	
ReVeil: Unconstrained Concealed Backdoor Attack on Deep Neural Networks Using Machine Unlearning	1046
<i>Manaar Alam, Hithem Lamri, Michail Maniatakos</i>	
CND-IDS: Continual Novelty Detection for Intrusion Detection Systems	1053
<i>Sean Fuhrman, Onat Gungor, Tajana Rosing</i>	
Graph in the Vault: Protecting Edge GNN Inference with Trusted Execution Environment.....	1060
<i>Ruyi Ding, Tianhong Xu, Aidong Adam Ding, Yunsi Fei</i>	
Ensembler: Protect Collaborative Inference Privacy from Model Inversion Attack Via Selective Ensemble	1067
<i>Dancheng Liu, Chenhui Xu, Jiajie Li, Amir Nassereldine, Jinjun Xiong</i>	
CAE-DFKD: Bridging the Transferability Gap in Data-Free Knowledge Distillation.....	1074
<i>Zherui Zhang, Changwei Wang, Rongtao Xu, Wenhao Xu, Shibiao Xu, Yu Zhang, Jie Zhou, Li Guo</i>	
On Bit-Level Reverse Engineering of Vehicular CAN Bus.....	1081
<i>Yunlang Cai, Hanxue Shi, Xiaohang Wang, Haoting Shen, Li Lu, Kui Ren</i>	
HoBBY: Hardening Unbalanced Branches Against Control Flow Attacks on Intel SGX and AMD SEV	1088
<i>Chang Liu, Shuaihu Feng, Yuan Li, Dongsheng Wang, Trevor E. Carlson</i>	
CMFuzz: Parallel Fuzzing of IoT Protocols by Configuration Model Identification and Scheduling	1095
<i>Qi Xu, Fuchen Ma, Yuanliang Chen, Wanli Chen, Feifan Wu, Yanyang Zhao, Heyuan Shi, Yu Jiang</i>	
A Novel Covert Timing Channel for Cloud FPGAs	1102
<i>Brian Udugama, Darshana Jayasinghe, Hassaan Saadat, Aleksandar Ignjatovic, Sri Parameswaran</i>	
RAP-Track: Efficient Control Flow Attestation Via Parallel Tracking in Commodity MCUs	1109
<i>Antonio Joia Neto, Adam Caulfield, Ivan De Oliveira Nunes</i>	
An Enhanced Data Packing Method for General Matrix Multiplication in Brakerski/Fan-Vercauteren Scheme	1116
<i>Xiangchen Meng, Yan Tan, Zijun Jiang, Yangdi Lyu</i>	
NSFlow: An End-To-End FPGA Framework with Scalable Dataflow Architecture for Neuro-Symbolic AI.....	1123
<i>Hanchen Yang, Zishen Wan, Ritik Raj, Joongun Park, Ziwei Li, Ananda Samajdar, Arijit Raychowdhury, Tushar Krishna</i>	
BitPattern: Enabling Efficient Bit-Serial Acceleration of Deep Neural Networks Through Bit-Pattern Pruning.....	1130
<i>Gang Wang, Siqi Cai, Zhenyu Li, Wenjie Li, Dongxu Lyu, Yanan Sun, Jianfei Jiang, Guanghui He</i>	
BLOOM: Bit-Slice Framework for DNN Acceleration with Mixed-Precision	1137
<i>Fangxin Liu, Ning Yang, Zongwu Wang, Xuanpeng Zhu, Haidong Yao, Xiankui Xiong, Li Jiang, Haibing Guan</i>	

ESM: A Framework for Building Effective Surrogate Models for Hardware-Aware Neural Architecture Search	1144
<i>Azaz-Ur-Rehman Nasir, Samroz Ahmad Shoaib, Muhammad Abdullah Hanif, Muhammad Shafique</i>	
SNAPPiX: Efficient-Coding-Inspired In-Sensor Compression for Edge Vision	1151
<i>Weikai Lin, Tianrui Ma, Adith Boloor, Yu Feng, Ruofan Xing, Xuan Zhang, Yuhao Zhu</i>	
Introducing Instruction-Accurate Simulators for Performance Estimation of Autotuning Workloads	1158
<i>Rebecca Pelke, Nils Bosbach, Lennart M. Reimann, Rainer Leupers</i>	
Speculative Decoding for Verilog: Speed and Quality, All in One	1165
<i>Changran Xu, Yi Liu, Yunhao Zhou, Shan Huang, Ningyi Xu, Qiang Xu</i>	
PARO: Hardware-Software Co-Design with Pattern-Aware Reorder-Based Attention Quantization in Video Generation Models	1172
<i>Xinhao Yang, Tianchen Zhao, Hongyi Wang, Wenheng Ma, Shulin Zeng, Zhenhua Zhu, Xuefei Ning, Huazhong Yang, Yu Wang</i>	
Hybrid Embedding Framework for Memory-Efficient Recommendation Systems	1179
<i>Seung Jin Yang, Hyuk-Jae Lee, Chae Eun Rhee</i>	
Mixed-Precision Quantization for Deep Vision Models with Integer Quadratic Programming	1186
<i>Zihao Deng, Sayeh Sharify, Xin Wang, Michael Orshansky</i>	
Maximizing Energy Efficiency in Spiking Neural Networks: A Dynamic Joint Pruning Framework	1193
<i>Shuo Chen, Zeshi Liu, Haihang You</i>	
DCDiff: Enhancing JPEG Compression Via Diffusion-Based DC Coefficients Estimation	1200
<i>Ziyuan Zhang, Han Qiu, Tianwei Zhang, Bin Chen, Chao Zhang</i>	
BiNeuroRAM: Energy-Efficient ReRAM-Based PIM for Accurate Bipolar Spiking Neural Network Acceleration.....	1207
<i>Jun Yan Lee, Chen Nie, Kang You, Yueyang Jia, Rui Yang, Zhezhi He</i>	
ResISC: Residue Number System-Based Integrated Sensing and Computing for Efficient Edge AI	1214
<i>Sepehr Tabrizchi, Samin Sohrabi, Mohamadreza Mohammadi, Ramtin Zand, Shaahin Angizi, Arman Roohi</i>	
Efficient Edge Vision Transformer Accelerator with Decoupled Chunk Attention and Hybrid Computing-In-Memory	1221
<i>Yi Li, Zijian Ye, Xiangqu Fu, Songqi Wang, Shucheng Du, Ning Lin, Dashan Shang, Jinshan Yue, Zhongrui Wang, Xiaojuan Qi, Feng Zhang, Han Wang</i>	
3D-CIMlet: A Chiplet Co-Design Framework for Heterogeneous In-Memory Acceleration of Edge LLM Inference and Continual Learning	1228
<i>Shuting Du, Luqi Zheng, Aradhana Mohan Parvathy, Feifan Xie, Tiwei Wei, Anand Raghunathan, Haitong Li</i>	
BlockPIM: Optimizing Memory Management for PIM-Enabled Long-Context LLM Inference	1235
<i>Zhichun Li, Jun Zhou, Xueqi Li, Ninghui Sun</i>	
PIMPAL: Accelerating LLM Inference on Edge Devices Via In-DRAM Arithmetic Lookup	1242
<i>Yoonho Jang, Hyeongjun Cho, Yesin Ryu, Junrae Kim, Seokin Hong</i>	

Free and Fair Hardware: A Pathway to Copyright Infringement-Free Verilog Generation Using LLMs.....	1249
<i>Sam Bush, Matthew Delorenzo, Phat Tieu, Jeyavijayan Rajendran</i>	
Hardware Generation with High Flexibility Using Reinforcement Learning Enhanced LLMs	1256
<i>Yifang Zhao, Weimin Fu, Shijie Li, Yi-Xiang Hu, Xiaolong Guo, Yier Jin</i>	
NetTAG: A Multimodal RTL-And-Layout-Aligned Netlist Foundation Model Via Text-Attributed Graph.....	1263
<i>Wenji Fang, Wenkai Li, Shang Liu, Yao Lu, Hongce Zhang, Zhiyao Xie</i>	
Enhancing LLM-Based Quantum Code Generation with Multi-Agent Optimization and Quantum Error Correction	1270
<i>Charlie Campbell, Hao Mark Chen, Wayne Luk, Hongxiang Fan</i>	
SynCircuit: Automated Generation of New Synthetic RTL Circuits Can Enable Big Data in Circuits.....	1277
<i>Shang Liu, Jing Wang, Wenji Fang, Zhiyao Xie</i>	
PyraNet: A Multi-Layered Hierarchical Dataset for Verilog	1284
<i>Bardia Nadimi, Ghali Omar Boutaib, Hao Zheng</i>	
PatLabor: Pareto Optimization of Timing-Driven Routing Trees.....	1291
<i>Zhiyang Chen, Hailong Yao, Xia Yin</i>	
Cost-Distance Steiner Trees for Timing-Constrained Global Routing	1298
<i>Stephan Held, Edgar Perner</i>	
Dynamic Local Usage: An Accurate Model for Usage of Tile-Internal Wiring in Global Routing.....	1304
<i>Tilman Bihler, Daniel Blankenburg</i>	
Reinforcement Learning-Driven Window Selection for Enhanced Window-Based Rip-Up and Reroute in Chip Detailed Routing	1310
<i>Yu-Chan Keng, Yu-Chun Pai, Wen-Hao Liu, Haoxing Ren, Danny Liu, Rongjian Liang, Mark Ho, Anthony Agnesina, Yih-Lang Li</i>	
Mr.TPL: A Method for Multi-Pin Net Router in Triple Patterning Lithography	1317
<i>Chengkai Wang, Weiqing Ji, Mingyang Kou, Zhiyang Chen, Fei Li, Nengyong Zhu, Hailong Yao</i>	
TransRoute: A Novel Hierarchical Transistor-Level Routing Framework Beyond Standard-Cell Methodology	1323
<i>Chen-Hao Hsu, David Z. Pan, Laurent Perron, Frédéric Didier, Xiaoqing Xu, Hao Chen</i>	
Asymmetric Predictive Testing for Aging in SRAMs	1330
<i>Yunkun Lin, Mingye Li, Sandeep Gupta</i>	
Machine Learning-Driven STL Generation for Enhancing Functional Safety of E/E Systems.....	1337
<i>Sanjay Das, Swastik Bhattacharya, Anand Menon, Shamik Kundu, Pooja Madhusoodhanan, Prasanth Viswanathan Pillai, Rubin Parekhji, Arnab Raha, Suvadeep Banerjee, Suriyaprakash Natarajan, Kanad Basu</i>	
EPICS: Efficient Parallel Pattern Fault Simulation for Sequential Circuits Via Strongly Connected Components.....	1344
<i>Mingjun Wang, Hui Wang, Jianan Mu, Xinyu Zhang, Bin Sun, Yihan Wen, Zizhen Liu, Feng Gu, Jun Gao, Shengwen Liang, Jing Ye, Xiaowei Li, Huawei Li</i>	

PastATPG: A Hybrid ATPG Framework for Better Test Compaction with Partial Assignment SAT	1351
<i>Zhiteng Chao, Xindi Zhang, Xinyu Zhang, Jianan Mu, Zizhen Liu, Shengwen Liang, Shaowei Cai, Jing Ye, Xiaowei Li, Huawei Li</i>	
Process-Variation-Aware Design Optimization for Wavelength-Routed Optical Networks-On-Chip.....	1358
<i>Liaoyuan Cheng, Mengchu Li, Tsun-Ming Tseng, Martin Schottenloher, Ulf Schlichtmann</i>	
FT-MUX: A Fault-Tolerant Microfluidic Multiplexer Design.....	1365
<i>Mengchu Li, Jiahui Peng, Tsun-Ming Tseng, Ulf Schlichtmann</i>	
VEDA: Efficient LLM Generation Through Voting-Based KV Cache Eviction and Dataflow-Flexible Accelerator.....	1372
<i>Zhican Wang, Hongxiang Fan, Haroon Waris, Gang Wang, Zhenyu Li, Jianfei Jiang, Yanan Sun, Guanghui He</i>	
LLMShare: Optimizing LLM Inference Serving with Hardware Architecture Exploration.....	1379
<i>Hongduo Liu, Chen Bai, Peng Xu, Lihao Yin, Xianzhi Yu, Hui-Ling Zhen, Mingxuan Yuan, Tsung-Yi Ho, Bei Yu</i>	
SSDTrain: An Activation Offloading Framework to SSDs for Faster Large Language Model Training	1386
<i>Kun Wu, Jeongmin Brian Park, Xiaofan Zhang, Mert Hidayetoglu, Vikram Sharma Mailthody, Sitao Huang, Steve Lumetta, Wen-Mei Hwu</i>	
LEMOE: LLM-Enhanced Multi-Objective Bayesian Optimization for Microarchitecture Exploration	1393
<i>Jingyuan Li, Jianrong Zhang, Ye Li, Wenbo Yin, Lingli Wang</i>	
SpecASR: Accelerating LLM-Based Automatic Speech Recognition Via Speculative Decoding	1400
<i>Linye Wei, Shuzhang Zhong, Songqiang Xu, Runsheng Wang, Ru Huang, Meng Li</i>	
PISA: Efficient Precision-Slice Framework for LLMs with Adaptive Numerical Type.....	1407
<i>Ning Yang, Zongwu Wang, Qingxiao Sun, Liqiang Lu, Fangxin Liu</i>	
SAPO: Improving the Scalability and Accuracy of Quantum Linear Solver for Portfolio Optimization.....	1414
<i>Tianze Zhu, Liqiang Lu, Jiajun Chen, Yuhang Chen, Hengrui Chen, Meng Xi, Jinshan Zhang, Xiaoming Sun, Jianwei Yin</i>	
Scalable Community Detection Using Quantum Hamiltonian Descent and QUBO Formulation.....	1421
<i>Jinglei Cheng, Ruilin Zhou, Yuhang Gan, Chen Qian, Junyu Liu</i>	
Computational Advantage in Hybrid Quantum Neural Networks: Myth Or Reality?	1428
<i>Muhammad Kashif, Alberto Marchisio, Muhammad Shafique</i>	
Measurement-Based Uncomputation of Quantum Circuits for Modular Arithmetic.....	1435
<i>Alessandro Luongo, Antonio Michele Miti, Varun Narasimhachar, Adithya Sireesh</i>	
Optimizing Windowed Arithmetic for Quantum Attacks Against RSA-2048	1442
<i>Alessandro Luongo, Varun Narasimhachar, Adithya Sireesh</i>	
Hardware-Software Co-Design for Distributed Quantum Computing	1449
<i>Ji Liu, Allen Zang, Martin Suchara, Tian Zhong, Paul D Hovland</i>	
Espresso: Exploiting the Sparsity Property in Event Sensors with Spatiotemporal Ordering	1455
<i>Leshan Li, Hongyi Li, Qingyuan Yang, Mingtao Ou, Rong Zhao, Xinglong Ji</i>	

CognitiveArm: Enabling Real-Time EEG-Controlled Prosthetic Arm Using Embodied Machine Learning	1462
<i>Abdul Basit, Maha Nawaz, Saim Rehman, Muhammad Shafique</i>	
Learning-Aided Safe Controller Synthesis with Formal Guarantees Via Vector Barrier Certificates	1469
<i>Xia Zeng, Mengxin Ren, Zhiming Liu, Zhengfeng Yang</i>	
Live Region Mutation Testing for Commercial Cyber-Physical System Development Tool Chain.....	1476
<i>Lehuan Zhang, Shikai Guo, Zixuan Wang, Xiaoyu Wang, Xiaochen Li, He Jiang</i>	
MAS-ISP: A Proxy-Free Online Hyperparameter Optimization Framework for ISP Hardware System	1483
<i>Jiaming Liu, Xuan Huang, Zhijian Hao, Ruoxi Zhu, Qi Zheng, Shuocheng Wang, Shushi Chen, Chang Liu, Leilei Huang, Jun Tao, Yibo Fan</i>	
NoiseZO: RRAM Noise-Driven Zeroth-Order Optimization for Efficient Forward-Only Training	1490
<i>Shuqi Wang, Zhengwu Liu, Chenchen Ding, Chen Zhang, Taiqiang Wu, Jiajun Zhou, Ngai Wong</i>	
APSQ: Additive Partial Sum Quantization with Algorithm-Hardware Co-Design	1497
<i>Yonghao Tan, Pingcheng Dong, Yongkun Wu, Yu Liu, Xuejiao Liu, Peng Luo, Shih-Yang Liu, Xijie Huang, Dong Zhang, Luhong Liang, Kwang-Ting Cheng</i>	
NN-AdderNet: Nonnegative and Sparse Weight Optimization Towards Ultra-Low Bitwidth AdderNet Quantization and Compression.....	1504
<i>Yunxiang Zhang, Gengchen Sun, Lizhi Fang, Biao Sun, Wenfeng Zhao</i>	
Replay4NCL: An Efficient Memory Replay-Based Methodology for Neuromorphic Continual Learning in Embedded AI Systems	1511
<i>Mishal Fatima Minhas, Rachmad Vidya Wicaksana Putra, Falah Awwad, Osman Hasan, Muhammad Shafique</i>	
FF-INT8: Efficient Forward-Forward DNN Training on Edge Devices with INT8 Precision	1518
<i>Jingxiao Ma, Priyadarshini Panda, Sherief Reda</i>	
BirdMoE: Reducing Communication Costs for Mixture-Of-Experts Training Using Load-Aware Bi-Random Quantization.....	1525
<i>Donglei Wu, Weihao Yang, Xiangyu Zou, Jinda Jia, Dingwen Tao, Wen Xia, Zhihong Tian</i>	
HeteroSVD: Efficient SVD Accelerator on Versal ACAP with Algorithm-Hardware Co-Design	1532
<i>Xinya Luan, Zhe Lin, Kai Shi, Jianwang Zhai, Kang Zhao</i>	
VSpGEMM: Exploiting Versal ACAP for High-Performance SpGEMM Acceleration.....	1539
<i>Kai Shi, Zhe Lin, Xinya Luan, Jianwang Zhai, Kang Zhao</i>	
HiSpTRSV: Exploring Tile-Level Parallelism for SpTRSV Acceleration on FPGAs.....	1546
<i>Fan Sun, Fang Dong, Dian Shen</i>	
A Data-Centric Hardware Accelerator for Efficient Adaptive Radix Tree.....	1553
<i>Jin Zhao, Yu Zhang, Jun Huang, Weihang Yin, Hui Yu, Hao Qi, Zixiao Wang, Longlong Lin, Xiaofei Liao, Hai Jin</i>	
ALLMod: Exploring Area-Efficiency of LUT-Based Large Number Modular Reduction Via Hybrid Workloads.....	1560
<i>Fangxin Liu, Haomin Li, Zongwu Wang, Bo Zhang, Mingzhe Zhang, Shoumeng Yan, Li Jiang, Haibing Guan</i>	

GPS: GNN-Based Two-Stage Pre-Scheduling Loop Mapping Method on CGRAs.....	1567
<i>Mingyang Kou, Weiqing Ji, Shouyi Yin, Hailong Yao</i>	
Rewire: Advancing CGRA Mapping Through a Consolidated Routing Paradigm.....	1574
<i>Zhaoying Li, Dan Wu, Dhananjaya Wijerathne, Dan Chen, Huize Li, Cheng Tan, Tulika Mitra</i>	
Routability-Aware Packing for High-Density Nonvolatile FPGAs.....	1581
<i>Huichuan Zheng, Yuqing Xiong, Jian Zuo, Hao Zhang, Zheng Jia, Mengying Zhao</i>	
EPIC: Error Prediction and Correction for Power-Efficient Voltage Underscaling Multiply-Accumulate Unit.....	1588
<i>Tongjing Wu, Xiaolu Hu, Tong Li, Siting Liu, Hui Wang, Weifeng He, Zhigang Mao, Honglan Jiang</i>	
PoP-ECC: Robust and Flexible Error Correction Against Multi-Bit Upsets in DNN Accelerators.....	1595
<i>Taewon Park, Saeid Gorgin, Dongwhae Kim, Jaeho Shin, Michael B. Sullivan, Jungrae Kim</i>	
AdreamDCO: AI-Driven Robust and Efficient Design Automation for Digitally Controlled Oscillators.....	1602
<i>Yaolong Hu, Hao Guo, Shikai Wang, Jiaqi Liu, Weidong Cao, Taiyun Chi</i>	
EVA: An Efficient and Versatile Generative Engine for Targeted Discovery of Novel Analog Circuits.....	1609
<i>Jian Gao, Weimin Fu, Xiaolong Guo, Weidong Cao, Xuan Zhang</i>	
BS-PDN-Last: Towards Optimal Power Delivery Network Design with Multifunctional Backside Metal Layers.....	1616
<i>Min Gyu Park, Amaan Rahman, Sung Kyu Lim</i>	
Towards Uncertainty-Aware Robotic Perception Via Mixed-Signal BNN Engine Leveraging Probabilistic Quantum Tunneling.....	1623
<i>Likai Pei, Yu Zhou, Xingtian Wang, Xueji Zhao, Wanxin Huang, Boyang Cheng, Halid Mulaosmanovic, Stefan Duenkel, Dominik Kleimaier, Sven Beyer, Kai Ni, Mengxue Hou, Michael Niemier, Ningyuan Cao</i>	
Dual-Issue Execution of Mixed Integer and Floating-Point Workloads on Energy-Efficient In-Order RISC-V Cores.....	1630
<i>Luca Colagrande, Luca Benini</i>	
A High-Precision and Low-Cost Approximate Transform Accelerator for Video Coding.....	1637
<i>Zhijian Hao, Jiaming Liu, Chenlong He, Qi Zheng, Shushi Chen, Jinchang Xu, Yue Hao, Xiao Yan, Xiaohua Ma</i>	
3D-TokSIM: Stacking 3D Memory with Token-Stationary Compute-In-Memory for Speculative LLM Inference.....	1644
<i>Wentao Zhao, Boya Lv, Meng Wu, Peiyu Chen, Fengyun Yan, Yufei Ma, Tianyu Jia, Ru Huang, Le Ye</i>	
A Memory-Efficient LLM Accelerator with Q-K Correlation Prediction Using Cluster-Based Associative Array for Selective KV Accessing.....	1651
<i>Zikang Zhou, Kaiqi Chen, Xuyang Duan, Jun Han</i>	
Precon: A Precision-Convertible Architecture for Accelerating Quantized Deep Learning Models Across Various Domains Including LLMs.....	1658
<i>Jongwoo Park, Hyeonseong Kim, Jiyun Han, Seungkyu Choi</i>	

RADiT: Redundancy-Aware Diffusion Transformer Acceleration Leveraging Timestep Similarity	1665
<i>Youngjun Park, Sangyeon Kim, Yeonggeon Kim, Gisan Ji, Sungju Ryu</i>	
SQ-DM: Accelerating Diffusion Models with Aggressive Quantization and Temporal Sparsity	1672
<i>Zichen Fan, Steve Dai, Rangharajan Venkatesan, Dennis Sylvester, Brucek Khailany</i>	
XShift: FPGA-Efficient Binarized LLM with Joint Quantization and Sparsification	1679
<i>Shuai Zhou, Huinan Tian, Sisi Meng, Jianli Chen, Jun Yu, Kun Wang</i>	
BBAL: A Bidirectional Block Floating Point-Based Quantisation Accelerator for Large Language Models	1686
<i>Xiaomeng Han, Yuan Cheng, Jing Wang, Junyang Lu, Hui Wang, X. X. Zhang, Ning Xu, Dawei Yang, Zhe Jiang</i>	
An Algorithm-Hardware Co-Design Based on Revised Microscaling Format Quantization for Accelerating Large Language Models	1693
<i>Yingbo Hao, Huangxu Chen, Yi Zou, Yanfeng Yang</i>	
Gem5-AcceSys: Enabling System-Level Exploration of Standard Interconnects for Novel Accelerators	1700
<i>Qunyou Liu, Marina Zapater, David Atienza</i>	
Adora Compiler: End-To-End Optimization for High-Efficiency Dataflow Acceleration and Task Pipelining on CGRAs	1707
<i>Jiahang Lou, Qilong Zhu, Yuan Dai, Zewei Zhong, Wenbo Yin, Lingli Wang</i>	
Automated Generation of Decoders for Irregular Instruction Sets Using Information-Theoretic Decision Tree Construction Algorithms	1714
<i>Lillian Tadros</i>	
Look Before You Leap: A Self-Review Bayesian Optimization Method for Constrained High-Dimensional Design Space Exploration	1721
<i>Xuyang Zhao, Yiyang Zhao, Zheng Wu, Tianning Gao, Zhaori Bi, Changhao Yan, Dian Zhou, Xuan Zeng</i>	
High-Performance Computing Architecture Exploration with Stage-Enhanced Bayesian Optimization	1728
<i>Vincent Fu, Mohamed Benazouz, Lilia Zaourar, Alix Munier-Kordon</i>	
On Design Space Exploration of Cache System in Multi-Chiplet Systems	1735
<i>Yan Zhang, Xiaohang Wang, Yingtao Jiang, Amit Kumar Singh</i>	
From Flatland to Forest: Exploring Pareto-Optimal Design Through RTL Hierarchy Trees	1742
<i>Donger Luo, Qi Sun, Xingheng Li, Cheng Zhuo, Bei Yu, Hao Geng</i>	
Synergistic Die-Level Router for Multi-FPGA System with Time-Division Multiplexing Optimization	1749
<i>Jiarui Wang, Yanjing Liu, Yibo Lin</i>	
BoolE: Exact Symbolic Reasoning Via Boolean Equality Saturation	1756
<i>Jiaqi Yin, Zhan Song, Chen Chen, Qihao Hu, Cunxi Yu</i>	
E-Morphic: Scalable Equality Saturation for Structural Exploration in Logic Synthesis	1763
<i>Chen Chen, Guangyu Hu, Cunxi Yu, Yuzhe Ma, Hongce Zhang</i>	
EDGE: DBMS-Empowered Boolean Decomposition for GIG Synthesis	1770
<i>Ruofei Tang, Xuliang Zhu, Xinyi Zhang, Lei Chen, Xing Li, Mingxuan Yuan, Jianliang Xu</i>	

Logic Restructuring with Preserved Logic Blocks	1777
<i>Siang-Yun Lee, Heinz Riener, Sascha Richter, Ankush Sood</i>	
ELF: Efficient Logic Synthesis by Pruning Redundancy in Refactoring	1783
<i>Dimitris Tsaras, Xing Li, Lei Chen, Zhiyao Xie, Mingxuan Yuan</i>	
Mixed Structural Choice Operator: Enhancing Technology Mapping with Heterogeneous Representations	1790
<i>Zhang Hu, Hongyang Pan, Yinshui Xia, Lunyao Wang, Zhufei Chu</i>	
Rank-Based Multi-Objective Approximate Logic Synthesis Via Monte Carlo Tree Search	1797
<i>Yuyang Ye, Xiangfei Hu, Yuchen Liu, Peng Xu, Yu Gong, Tinghuan Chen, Hao Yan, Bei Yu, Longxing Shi</i>	
Harrow: Synthesis of Optical Logic Circuits Via Harmonic Mean and Integer Partition.....	1804
<i>Jun-Wei Liang, Iris Hui-Ru Jiang, Kai-Hsiang Chiu</i>	
SSpMV: A Sparsity-Aware SpMV Framework Empowered by Multimodal Machine Learning	1811
<i>Shengle Lin, Chubo Liu, Yan Ding, Joey Tianyi Zhou, Kenli Li, Wangdong Yang</i>	
An Input-Aware Sparse Tensor Compiler Empowered by Vectorized Acceleration.....	1818
<i>Xianhao He, Haotian Wang, Jiapeng Zhang, Wangdong Yang, Anthony Theodore Chronopoulos, Kenli Li</i>	
Enabling On-Tiny-Device Model Personalization Via Gradient Condensing and Alternant Partial Update	1825
<i>Zhenge Jia, Yiyang Shi, Zeyu Bao, Zirui Wang, Xin Pang, Huiguo Liu, Yu Duan, Zhaoyan Shen, Mengying Zhao</i>	
Unlocking a New Rust Programming Experience: Fast and Slow Thinking with LLMs to Conquer Undefined Behaviors	1832
<i>Renshuang Jiang, Pan Dong, Zhenling Duan, Yu Shi, Xiaoxiang Fang, Yan Ding, Jun Ma, Shuai Zhao, Zhe Jiang</i>	
DROIDFUZZ: Proprietary Driver Fuzzing for Embedded Android Devices	1839
<i>Jianzhong Liu, Yuheng Shen, Yifei Chu, Qiang Zhang, Heyuan Shi, Wanli Chang, Yu Jiang</i>	
STREAM: Spatiotemporal Similarity-Based Efficient Approximate Median with Tunable Granularity.....	1846
<i>Fenfang Li, Huizhang Luo, Weichen Liu, Anthony Theodore Chronopoulos, Kenli Li, Chubo Liu</i>	
Enabling Data-Deduplication-Assisted Data Relocation for Interlaced Magnetic Recording.....	1852
<i>Chen-Jui Tu, Shuo-Han Chen</i>	
Location is Key: Leveraging LLM for Functional Bug Localization in Verilog Design	1858
<i>Bingkun Yao, Ning Wang, Jie Zhou, Xi Wang, Hong Gao, Zhe Jiang, Nan Guan</i>	
PICK: An SRAM-Based Processing-In-Memory Accelerator for K-Nearest-Neighbor Search in Point Clouds	1865
<i>Chen Nie, Chao Jiang, Liming Xiao, Weifeng Zhang, Zhezhi He</i>	
HH-PIM: Dynamic Optimization of Power and Performance with Heterogeneous-Hybrid PIM for Edge AI Devices	1872
<i>Sangmin Jeon, Kangju Lee, Kyeongwon Lee, Woojoo Lee</i>	

Anchor First, Accelerate Next: Revolutionizing GNNs with PIM by Harnessing Stationary Data.....	1879
<i>Jiaxian Chen, Yuxuan Qi, Yongbiao Zhu, Jianan Yuan, Kaoyi Sun, Tianyu Wang, Chenlin Ma, Yi Wang</i>	
3D-SubG: A 3D Stacked Hybrid Processing Near/In-Memory Accelerator for Subgraph GNNs.....	1886
<i>Guoxiang Li, Runnan Xu, Ruohang Xu, Yikan Qiu, Renati Tuerhong, Muhan Zhang, Le Ye, Yufei Ma</i>	
PIMDup: An Optimized Deduplication Design on a Real Processing-In-Memory System	1893
<i>Chun-Le Yeh, Liang-Chi Chen, Chien-Chung Ho, Yu-Ming Chang, Da-Wei Chang</i>	
An Efficient Compute-In-Memory Based Accelerator for Point-Based Point Cloud Neural Networks	1900
<i>Xipeng Lin, Cong Wang, Shanshi Huang, Hongwu Jiang</i>	
NDFT: Accelerating Density Functional Theory Calculations Via Hardware/Software Co-Design on Near-Data Computing System.....	1907
<i>Qingcai Jiang, Buxin Tu, Xiaoyu Hao, Junshi Chen, Hong An</i>	
CIMFlow: An Integrated Framework for Systematic Design and Evaluation of Digital CIM Architectures	1914
<i>Yingjie Qi, Jianlei Yang, You Wang, Yikun Wang, Dayu Wang, Ling Tang, Cenlin Duan, Xiaolin He, Weisheng Zhao</i>	
ChatLS: Multimodal Retrieval-Augmented Generation and Chain-Of-Thought for Logic Synthesis Script Customization	1921
<i>Haisheng Zheng, Haoyuan Wu, Zhuolun He</i>	
MAGE: A Multi-Agent Engine for Automated RTL Code Generation	1928
<i>Yujie Zhao, Hejia Zhang, Hanxian Huang, Zhongming Yu, Jishen Zhao</i>	
ReChisel: Effective Automatic Chisel Code Generation by LLM with Reflection	1935
<i>Juxin Niu, Xiangfeng Liu, Dan Niu, Xi Wang, Zhe Jiang, Nan Guan</i>	
Efficient Continuous Logic Optimization with Diffusion Model.....	1942
<i>Yikang Ouyang, Xiaofei Yu, Jiadong Zhu, Tinghuan Chen, Yuzhe Ma</i>	
SmaRTLy: RTL Optimization with Logic Inferencing and Structural Rebuilding	1949
<i>Chengxi Li, Yang Sun, Lei Chen, Yiwen Wang, Mingxuan Yuan, Evangeline F. Y. Young</i>	
Centralized Training and Decentralized Control Through the Actor-Critic Paradigm for Highly Optimized Multicores.....	1956
<i>Benedikt Dietrich, Heba Khdr, Jorg Henkel</i>	
Few-Shot Learning on AMS Circuits and Its Application to Parasitic Capacitance Prediction.....	1963
<i>Shan Shen, Yibin Zhang, Hector Rodriguez Rodriguez, Wenjian Yu</i>	
Decoupling Analog Circuit Representation from Technology for Behavior-Centric Optimization.....	1970
<i>Jintao Li, Haochang Zhi, Jiang Xiao, Keren Zhu, Yun Li</i>	
MOSS: Multi-Modal Representation Learning on Sequential Circuits	1977
<i>Mingjun Wang, Bin Sun, Jianan Mu, Feng Gu, Boyu Han, Tianmeng Yang, Xinyu Zhang, Silin Liu, Yihan Wen, Hui Wang, Jun Gao, Zhiteng Chao, Husheng Han, Zizhen Liu, Shengwen Liang, Jing Ye, Bei Yu, Xiaowei Li, Huawei Li</i>	
InsightAlign: A Transferable Physical Design Recipe Recommender Based on Design Insights	1984
<i>Hao-Hsiang Hsiao, Sudipto Kundu, Wei Zeng, Wei-Ting J. Chan, Deyuan Guo, Sung Kyu Lim</i>	

Self-Attention to Operator Learning-Based 3D-IC Thermal Simulation.....	1991
<i>Zhen Huang, Hong Wang, Wenkai Yang, Muxi Tang, Depeng Xie, Ting-Jung Lin, Yu Zhang, Wei W. Xing, Lei He</i>	
Accelerating Design-Technology Co-Development Using Neural Compact Modeling and Data-Driven SPICE Simulation.....	1998
<i>Yongjeong Lee, Seungsoo Lee, Jeongyeol Kim, Jungyun Choi, Zhaojie Li, Dehuang Wu, Joddy Wang</i>	
PracMHBench: Re-Evaluating Model-Heterogeneous Federated Learning Based on Practical Edge Device Constraints	2004
<i>Yuanchun Guo, Bingyan Liu, Yulong Sha, Zhensheng Xian</i>	
SimPhony: A Device-Circuit-Architecture Cross-Layer Modeling and Simulation Framework for Heterogeneous Electronic-Photonic AI System.....	2011
<i>Ziang Yin, Meng Zhang, Nicholas Gangi, Rena Huang, Jeff Zhang, Jiaqi Gu</i>	
DataMaestro: A Versatile and Efficient Data Streaming Engine Bringing Decoupled Memory Access to Dataflow Accelerators	2018
<i>Xiaoling Yi, Yunhao Deng, Ryan Antonio, Fanchen Kong, Guilherme Paim, Marian Verhelst</i>	
ITaskSense: Task-Oriented Object Detection in Resource-Constrained Environments.....	2025
<i>Suncheon Jeong, Hamza Errahmouni Barkam, Hyunwoo Oh, Hanning Chen, Tamoghno Das, Zhen Ye, Mohsen Imani</i>	
PASK: Cold Start Mitigation for Inference with Proactive and Selective Kernel Loading on GPUs.....	2032
<i>Xuanteng Huang, Jiansu Du, Nong Xiao, Xianwei Zhang</i>	
FLAG: An FPGA-Based System for Low-Latency GNN Inference Service Using Vector Quantization	2039
<i>Yunki Han, Taehwan Kim, Jiwan Kim, Seohye Ha, Lee-Sup Kim</i>	
DSPlacer: DSP Placement for FPGA-Based CNN Accelerator.....	2046
<i>Baohui Xie, Xinrui Zhu, Zhiyuan Lu, Yuan Pu, Tongkai Wu, Xiaofeng Zou, Bei Yu, Tinghuan Chen</i>	
ReMaP: Macro Placement by Recursively Prototyping and Periphery-Guided Relocating.....	2053
<i>Yunqi Shi, Xi Lin, Siyuan Xu, Shixiong Kai, Ke Xue, Mingxuan Yuan, Chao Qian, Zhi-Hua Zhou</i>	
RUPlace: Optimizing Routability Via Unified Placement and Routing Formulation.....	2060
<i>Yifan Chen, Jing Mai, Zuodong Zhang, Yibo Lin</i>	
Differentiable Net-Moving and Local Congestion Mitigation for Routability-Driven Global Placement	2067
<i>Wenchao Li, Hongxi Wu, Duanxiang Liu, Xingquan Li, Wenxing Zhu</i>	
Comprehensive Placement and Routing Framework with Guaranteed In-Cell Routability for Synthesizing Complementary-FET Cells	2074
<i>Zhengzhe Zheng, Yinuo Wu, Keyu Peng, Chao Wang, Ziran Zhu</i>	
Synthesis of CFET Cell Library Leveraging Backside Metal Routing.....	2080
<i>Ting-Xin Lin, Yih-Lang Li</i>	
SuperCopyback: Revisiting Copyback on Modern High-Performance NAND Flash-Based SSDs	2087
<i>Dong Huang, Bo Ding, Wei Tong, Dan Feng</i>	

The Unwritten Contract of Cloud-Based Elastic Solid-State Drives.....	2094
<i>Yingjia Wang, Ming-Chang Yang</i>	
MiniWear: Minimizing Flash Wear Via Hybrid Persistent Cache for Extended EF-SMR Lifetime.....	2101
<i>Chenlin Ma, Kaoyi Sun, Yuxuan Qi, Jiaxian Chen, Xiaochuan Zheng, Tianyu Wang, Yi Wang</i>	
Leopard: Hardware Pass-Through Remote Storage Access with Queue Concurrency for Edge Intelligent Workstations.....	2108
<i>Wenjie Wang, Bo Peng, Jianguo Yao, Haibing Guan</i>	
FineRR-ZNS: Enabling Fine-Granularity Read Refreshing for ZNS SSDs.....	2115
<i>Jun Li, Zhibing Sha, Fan Yang, Xiaofei Xu, Xiaobai Chen, Jieming Yin, Jianwei Liao</i>	
StreamCSD: SSD-Autonomous Stream Management Via In-Storage Content Learning.....	2122
<i>Wenjie Li, Xiang Chen, Yelin Shan, Jiapin Wang, Yunxin Huang, Yafei Yang, Tao Lu, You Zhou, Fei Wu</i>	
Exploiting Power Side-Channel Vulnerabilities in XGBoost Accelerator.....	2129
<i>Yimeng Xiao, Archit Gajjar, Aydin Aysu, Paul Franzon</i>	
ZkVC: Fast Zero-Knowledge Proof for Private and Verifiable Computing.....	2136
<i>Yancheng Zhang, Mengxin Zheng, Xun Chen, Jingtong Hu, Weidong Shi, Lei Ju, Yan Solihin, Qian Lou</i>	
ABC-FHE: A Resource-Efficient Accelerator Enabling Bootstrappable Parameters for Client-Side Fully Homomorphic Encryption.....	2143
<i>S. Yune, H. Lee, A. Putra, H. Cho, C. Manh, J. Jeon, Y. Kim</i>	
SeDA: Secure and Efficient DNN Accelerators with Hardware/Software Synergy.....	2150
<i>Wei Xuan, Zhongrui Wang, Lang Feng, Ning Lin, Zihao Xuan, Rongliang Fu, Tsung-Yi Ho, Yuzhong Jiao, Luhong Liang</i>	
Guarder: A Stable and Lightweight Reconfigurable RRAM-Based PIM Accelerator for DNN IP Protection.....	2157
<i>Ning Lin, Yi Li, Jiankun Li, Jichang Yang, Yangu He, Yukui Luo, Dashan Shang, Xiaoming Chen, Xiaojuan Qi, Zhongrui Wang</i>	
Quorum: Zero-Training Unsupervised Anomaly Detection Using Quantum Autoencoders.....	2164
<i>Jason Zev Ludmir, Sophia Rebello, Jacob Ruiz, Tirthak Patel</i>	
Mitigating Routability Problems in Complementary-FET-Based VLSI Designs.....	2171
<i>Junghyun Yoon, Heechun Park</i>	
Design and Technology Co-Optimization Utilizing Flip-FET (FFET) Standard Cells.....	2178
<i>Jaehoon Ahn, Taewhan Kim</i>	
Multi-Agent Yield Analysis for Circuit Design.....	2185
<i>Haiyan Qin, Jing Kou, Liang Zhang, Wang Kang, Wei W. Xing</i>	
YAP: Yield Modeling and Simulation for Advanced Packaging.....	2192
<i>Zhichao Chen, Puneet Gupta</i>	
SDM-PEB: Spatial-Depthwise Mamba for Enhanced Post-Exposure Bake Simulation.....	2199
<i>Ziyang Yu, Peng Xu, Zixiao Wang, Binwu Zhu, Qipan Wang, Yibo Lin, Runsheng Wang, Bei Yu, Martin Wong</i>	

ChipletEM: Physics-Based 2.5D and 3D Chiplet Heterogeneous Integration Electromigration Signoff Tool Using Coupled Stress and Thermal Simulation	2206
<i>Zeyu Sun, Weijie Tong, Xiaoning Ma, He Cao, Jianyun Liu, Zhiqiang Li, Qinzi Xu</i>	
MHDiff: Memory- And Hardware-Efficient Diffusion Acceleration Via Focal Pixel Aware Quantization	2213
<i>Chunyu Qi, Xuhang Wang, Ruiyang Chen, Yuanzheng Yao, Naifeng Jing, Chen Zhang, Jun Wang, Zhihui Fu, Xiaoyao Liang, Zhuoran Song</i>	
Harnessing Conventional Video Processing Insights for Emerging 3D Video Generation Models: A Comprehensive Attention-Aware Way.....	2220
<i>Tianlang Zhao, Jun Liu, Xingyang Li, Li Ding, Jinhao Li, Shuaiheng Li, Jinbo Hu, Guohao Dai</i>	
STREAMINGGS: Voxel-Based Streaming 3D Gaussian Splatting with Memory Optimization and Architectural Support	2227
<i>Chenqi Zhang, Yu Feng, Jieru Zhao, Guangda Liu, Wenchao Ding, Chentao Wu, Minyi Guo</i>	
GauRast: Enhancing GPU Triangle Rasterizers to Accelerate 3D Gaussian Splatting.....	2234
<i>Sixu Li, Ben Keller, Yingyan Celine Lin, Brucek Khailany</i>	
Local-GS: An Order-Independent Gaussian Splatting Training Accelerator Exploiting Splat Locality.....	2241
<i>Yiyang Sun, Qinzhe Zhi, Yiqi Jing, Le Ye, Ru Huang, Tianyu Jia</i>	
CHORD: Composable Hybrid Optical Reconfigurable Diffractive Framework for Optical Neural Network.....	2248
<i>Ziang Yin, Yu Yao, Jeff Zhang, Jiaqi Gu</i>	
PiSPICE: Accelerating Post-Layout SPICE Simulation Via Essential Parasitic Identification	2255
<i>Zhou Jin, Jing Li, Jian Xin, Tianjia Zhou, Xiao Wu, Dan Niu, Zuochang Ye</i>	
Me-MPK: Accelerating Krylov Subspace Solvers Via Memory-Efficient Matrix-Power Kernel	2262
<i>Haozhong Qiu, Chuanfu Xu, Jianbin Fang, Shengguo Li, Liang Deng, Jian Zhang, Zhe Dai, Yue Ding, Yue Wang, Zhimeng Han, Yonggang Che, Jie Liu</i>	
New Time-Domain Preconditioners for HB Jacobian of RF Circuits	2269
<i>Chenyi Tan, Yangfeng Su, Fan Yang, Xuan Zeng</i>	
Efficient Recycling Subspace Truncation Method for Periodic Small-Signal Analysis	2276
<i>Yuncheng Xu, Fan Yang, Yangfeng Su</i>	
MemSens: Significantly Reducing Memory Overhead in Adjoint Sensitivity Analysis Using Novel Error-Bounded Lossy Compression	2283
<i>Chenxi Li, Yihang Feng, Fuxing Deng, Dingwen Tao, Weifeng Liu, Zhou Jin</i>	
MARIO: A Superadditive Multi-Algorithm Interworking Optimization Framework for Analog Circuit Sizing.....	2290
<i>Wangzhen Li, Yuan Meng, Ruiyu Lyu, Changhao Yan, Keren Zhu, Zhaori Bi, Dian Zhou, Xuan Zeng</i>	
PipeLink: A Pipelined Resource Sharing System for Dataflow High-Level Synthesis	2297
<i>Rui Li, Lincoln Berkley, Rajit Manohar</i>	
Optimizing Recovery Logic in Speculative High-Level Synthesis	2304
<i>Dylan Leothaud, Jean-Michel Gorius, Simon Rokicki, Steven Derrien</i>	

AutoClock: Automated Clock Management for Power-Efficient HLS Designs on FPGAs	2311
<i>Jiawei Liang, Linfeng Du, Xiaofeng Zhou, Zhe Lin, Jiang Xu, Wei Zhang</i>	
Cayman: Custom Accelerator Generation with Control Flow and Data Access Optimization.....	2318
<i>Youwei Xiao, Fan Cui, Zizhang Luo, Weijie Peng, Yun Liang</i>	
ADVISOR: Approximate Computing-FriendlY High-LeVel Synthesis DesIgn Space ExpLOrEr	2325
<i>B. Parchamdar, B. Carrion Schaefer</i>	
Comparison-Free Bit-Stream Generation for Cost-Efficient Unary Computing	2332
<i>Faeze S. Banitaba, Amir Hossein Jalilvand, M. Hassan Najafi, Sercan Aygun</i>	
3D-Flow: Flow-Based Standard Cell Legalization for 3D ICs.....	2339
<i>Yuxuan Zhao, Peiyu Liao, Bei Yu</i>	
DCO-3D: Differentiable Congestion Optimization in 3D ICs.....	2346
<i>Hao-Hsiang Hsiao, Yi-Chen Lu, Pruek Vanna-Iampikul, Anthony Agnesina, Rongjian Liang, Yuan-Hsiang Lu, Haoxing Ren, Sung Kyu Lim</i>	
GNN-MLS: Signal Routing in Mixed-Node 3D ICs Through GNN-Assisted Metal Layer Sharing	2353
<i>Jiawei Hu, Pruek Vanna-Iampikul, Zhen Zhuang, Tsung-Yi Ho, Sung Kyu Lim</i>	
A Systematic Approach for Multi-Objective Double-Side Clock Tree Synthesis	2360
<i>Xun Jiang, Haoran Lu, Yuxuan Zhao, Jiarui Wang, Zizheng Guo, Heng Wu, Bei Yu, Sung Kyu Lim, Runsheng Wang, Ru Huang, Yibo Lin</i>	
To Tackle Cost-Skew Tradeoff: An Adaptive Learning Approach for Hub Node Selection.....	2367
<i>Guowei Sun, Lin Chen, Qiming Huang, Hu Ding</i>	
FedEDA: Federated Learning Framework for Privacy-Preserving Machine Learning in EDA	2374
<i>Joonseok Kim, Donggyu Kim, Seonghyeon Park, Seokhyeong Kang</i>	
Ares: High Performance Near-Storage Accelerator for FHE-Based Private Set Intersection.....	2381
<i>Haoxuan Wang, Yinghao Yang, Jinkai Zhang, Hang Lu, Xiaowei Li</i>	
GraphAccel: An In-Storage Accelerator for Efficient Graph-Based Vector Similarity Search Using Page Packing and Speculative Search Optimization	2387
<i>Yoonyoung Kwon, Yunjong Boo, Hyungmin Cho</i>	
CIM-BLAS: Computing-In-Memory Accelerator for BLAS	2394
<i>Rui Liu, Zerun Li, Xiaoyu Zhang, Xiaoming Chen, Yinhe Han, Minghua Tang</i>	
Segmented Angular Pre-Processing for Accurate and Efficient In-Memory Vector Similarity Search	2401
<i>Chi-Tse Huang, Jen-Chieh Wang, Hsiang-Yun Cheng, An-Yeu Andy Wu</i>	
Efficient Weight Mapping and Resource Scheduling on Crossbar-Based Multi-Core CIM Systems.....	2408
<i>Hanjie Liu, Sifan Sun, Aifei Zhang, Haiyan Qin, Yutong Wu, Minhao Gu, Shihang Fu, Shuaikai Liu, Baosen Liu, Wang Kang</i>	
ARCANE: Adaptive RISC-V Cache Architecture for Near-Memory Extensions	2414
<i>Vincenzo Petrolo, Flavia Guella, Michele Caon, Pasquale Davide Schiavone, Guido Masera, Maurizio Martina</i>	
AcclMT: A Highly Resource-Efficient and Flexible Poseidon Hash-Based Merkle Tree Architecture.....	2420
<i>Changxu Liu, Hao Zhou, Lan Yang, Yifei Feng, Zheng Wu, Zhuoyuan Yang, Yinlong Li, Shiyong Wu, Fan Yang</i>	

LeakyDSP: Exploiting Digital Signal Processing Blocks to Sense Voltage Fluctuations in FPGAs.....	2427
<i>Xin Zhang, Jiajun Zou, Yi Yang, Qingni Shen, Zhi Zhang, Yansong Gao, Zhonghai Wu, Trevor E. Carlson</i>	
FastPath: A Hybrid Approach for Efficient Hardware Security Verification	2434
<i>Lucas Deutschmann, Andres Meza, Dominik Stoffel, Wolfgang Kunz, Ryan Kastner Rptu</i>	
Re4PUF: A Reliable, Reconfigurable ReRAM-Based PUF Resilient to DNN and Side Channel Attacks.....	2441
<i>Ning Lin, Yi Li, Yangu He, Songqi Wang, Hegan Chen, Kwunhang Wong, Chuxin Li, Jichang Yang, Yifei Yu, Meng Xu, Yongkang Han, Rui Chen, Xiaoming Chen, Xiaoxin Xu, Jianguo Yang, Dashan Shang, Zhongrui Wang</i>	
ACIM-QMM: Efficient Analog Computing-In-Memory Accelerator for QC-MDPC McEliece Cryptosystem.....	2448
<i>Pingdan Xiao, Zhengmiao Wei, Sichun Du, Wanli Chang, Qinghui Hong</i>	
AutoSkewBMT: Autonomously Synthesizing Optimized Integrity Authentication Mechanism for DNN Accelerators	2454
<i>Rakin Muhammad Shadab, Sanjay Gandham, Mingjie Lin</i>	
Hydra: Harnessing Expert Popularity for Efficient Mixture-Of-Expert Inference on Chiplet System.....	2461
<i>Siqi He, Haozhe Zhu, Jiawei Zheng, Lizhou Wu, Bo Jiao, Qi Liu, Xiaoyang Zeng, Chixiao Chen</i>	
SynGPU: Synergizing CUDA and Bit-Serial Tensor Cores for Vision Transformer Acceleration on GPU.....	2468
<i>Yuanzheng Yao, Chen Zhang, Chunyu Qi, Ruiyang Chen, Jun Wang, Zhihui Fu, Naifeng Jing, Xiaoyao Liang, Zhuoran Song</i>	
Buffer Prospector: Discovering and Exploiting Untapped Buffer Resources in Many-Core DNN Accelerators.....	2475
<i>Yuchen Wei, Jingwei Cai, Mingyu Gao, Sen Peng, Zuocong Wu, Guiming Shi, Kaisheng Ma</i>	
PacQ: A SIMT Microarchitecture for Efficient Dataflow in Hyper-Asymmetric GEMMs.....	2482
<i>Ruokai Yin, Yuhang Li, Priyadarshini Panda</i>	
MetaDSE: A Few-Shot Meta-Learning Framework for Cross-Workload CPU Design Space Exploration.....	2489
<i>Runzhen Xue, Hao Wu, Mingyu Yan, Ziheng Xiao, Xiaochun Ye, Dongrui Fan</i>	
SAGA: A Memory-Efficient Accelerator for GANN Construction Via Harnessing Vertex Similarity	2496
<i>Ruiyang Chen, Xueyuan Liu, Chunyu Qi, Yuanzheng Yao, Yanan Sun, Xiaoyao Liang, Zhuoran Song</i>	
DIAS: Distance-Based Attention Sparsity for Ultra-Long-Sequence Transformer with Tree-Like Processing-In-Memory Architecture	2503
<i>Zekai Chen, Yiming Chen, Teng Wan, Tianyi Yu, Yu Wang, Huazhong Yang, Xueqing Li</i>	
AttenPIM: Accelerating LLM Attention with Dual-Mode GEMV in Processing-In-Memory.....	2510
<i>Liyan Chen, Dongxu Lyu, Zhenyu Li, Jianfei Jiang, Qin Wang, Zhigang Mao, Naifeng Jing</i>	
SplitSync: Bank Group-Level Split-Synchronization for High-Performance DRAM PIM	2517
<i>Byungkuk Yoon, Sanghyeok Han, Gyeonghwan Park, Jae-Joon Kim</i>	

PIMoE: Towards Efficient MoE Transformer Deployment on NPU-PIM System Through Throttle-Aware Task Offloading.....	2524
<i>Lizhou Wu, Haozhe Zhu, Siqi He, Xuanda Lin, Xiaoyang Zeng, Chixiao Chen</i>	
Supporting Register-Based Addressing Modes for in-DRAM PIM ISAs	2531
<i>Seok Young Kim, Byung Ho Choi, Seokwon Kang, Yongjun Park, Seon Wook Kim</i>	
OutlierCIM: Outlier-Aware Digital CIM-Based LLM Accelerator with Hybrid-Strategy Quantization and Unified FP-INT Computation	2537
<i>Zihan Zou, Shikuang Chen, Chen Zhang, Xing Wang, Zhichao Liu, Haoran Du, Xin Si, Hao Cai, Bo Liu</i>	
Near-Memory LLM Inference Processor Based on 3D DRAM-To-Logic Hybrid Bonding	2544
<i>Sanghyeok Han, Byungkuk Yoon, Gyeonghwan Park, Choungki Song, Dongkyun Kim, Jae-Joon Kim</i>	
SeIM: In-Memory Acceleration for Approximate Nearest Neighbor Search	2551
<i>Chaoqiang Liu, Dan Chen, Yu Huang, Wenjing Xiao, Haifeng Liu, Yi Zhang, Huize Li, Xiaofei Liao, Hai Jin</i>	
Efficient Rectification Signal Validation for Optimal Functional ECO Patch Generation	2558
<i>Tzu-Yu Tung, Yu-Ling Hsu, Shao-Lun Huang, Chung-Yang Ric Huang</i>	
X-SAT: An Efficient Circuit-Based SAT Solver.....	2564
<i>Yuhang Qian, Zhihan Chen, Xindi Zhang, Shaowei Cai</i>	
Approximate SMT Counting Beyond Discrete Domains	2571
<i>Arijit Shaw, Kuldeep S. Meel</i>	
Leveraging Critical Proof Obligations for Efficient IC3 Verification.....	2578
<i>Lingfeng Zhu, Xindi Zhang, Yongjian Li, Shaowei Cai</i>	
Property-Driven Parallel Symbolic Model Checking of LTL.....	2585
<i>Yuheng Su, Yingcheng Li, Qiusong Yang, Yiwei Ci, Ziyu Huang</i>	
RE3: Finding Refinement Relations with Relational Mapping Abstraction.....	2592
<i>You Li, Guannan Zhao, Yunqi He, Hai Zhou</i>	
Logic Optimization Meets SAT: A Novel Framework for Circuit-SAT Solving	2599
<i>Zhengyuan Shi, Tiebing Tang, Jiaying Zhu, Sadaf Khan, Hui-Ling Zhen, Mingxuan Yuan, Zhufei Chu, Qiang Xu</i>	
H3Match: A Hybrid Heterogeneous Hypergraph Matching Method for Subcircuit Identification.....	2606
<i>Bohao Li, Qingsong Peng, Changhong Wang, Tianming Ni, Tinghuan Chen, Qi Sun, Cheng Zhuo</i>	
DAWN: Accelerating Point Cloud Object Detection Via Object-Aware Partitioning and 3D Similarity-Based Filtering	2613
<i>Dongdong Tang, Yu Mao, Weilan Wang, Nan Guan, Tei-Wei Kuo, Chun Jason Xue</i>	
Easz: An Agile Transformer-Based Image Compression Framework for Resource-Constrained IoTs	2620
<i>Yu Mao, Jingzong Li, Jun Wang, Hong Xu, Tei-Wei Kuo, Nan Guan, Chun Jason Xue</i>	
MMDFL: Multi-Model-Based Decentralized Federated Learning for Resource-Constrained AIoT Systems.....	2627
<i>Dengke Yan, Yanxin Yang, Ming Hu, Xin Fu, Mingsong Chen</i>	

LightRIM: Light Runtime Integrity Measurement for Linux Kernels in Embedded Applications.....	2634
<i>Yili Guo, Zhuoran Ma, Xiangyue Li, Jiajia Huang, Wanli Chang</i>	
Insights from Rights and Wrongs: A Large Language Model for Solving Assertion Failures in RTL Design.....	2640
<i>Jie Zhou, Youshu Ji, Ning Wang, Yuchen Hu, Xinyao Jiao, Bingkun Yao, Xinwei Fang, Shuai Zhao, Nan Guan, Zhe Jiang</i>	
A Post-Implementation Performance Prediction Method with HLS Optimization Directives	2647
<i>Jingyu Zhu, Yan Ding, Lu Xiao, Kenli Li, Chubo Liu, Zheng Xiao</i>	
Contention-Aware Forecasting of Energy Efficiency Through Sequence-Based Models in Modern Heterogeneous Processors	2653
<i>Mohammed Bakr Sikal, Jeferson González-Gómez, Heba Khdr, Jörg Henkel</i>	
AutoPower: Automated Few-Shot Architecture-Level Power Modeling by Power Group Decoupling	2660
<i>Qijun Zhang, Yao Lu, Mengming Li, Zhiyao Xie</i>	
MemSeer: Leverage Memory Failure Distinctions and Multi-Grained Prediction in Ultra-Scale Heterogeneous X86/ARM Clusters	2667
<i>Yunfei Gu, Yixuan Liu, Xinyuan Wu, Bo Shao, Chentao Wu, Shiyi Li, Jieru Zhao, Jie Li, Minyi Guo, Kunlin Yang, Wengui Zhang, Feilong Lin</i>	
CXL-ECC: An Efficient LRC-Based on-CXL-Memory-EXpander-Controller ECC to Enhance Reliability and Performance of DRAM Error Correction.....	2674
<i>Yixuan Liu, Yunfei Gu, Junhao Dai, Xinyuan Wu, Chentao Wu, Xinfei Guo, Jieru Zhao, Jie Li, Minyi Guo</i>	
Megabits Down to Kilobits: Memory-Efficient Time-Aware Shaping for TSN.....	2681
<i>Xuyan Jiang, Wenwen Fu, Xiangrui Yang, Yingwen Chen, Wenfei Wu, Zhigang Sun</i>	
AARC: Automated Affinity-Aware Resource Configuration for Serverless Workflows.....	2688
<i>Lingxiao Jin, Zinuo Cai, Zebin Chen, Hongyu Zhao, Ruhui Ma</i>	
FlexStep: Enabling Flexible Error Detection in Multi/Many-Core Real-Time Systems	2695
<i>Tinglue Wang, Yiming Li, Wei Tang, Jiapeng Guan, Zhenghui Guo, Renshuang Jiang, Ran Wei, Jing Li, Zhe Jiang</i>	
GraphFI: An Efficient Fault Injection Framework for Graph Processing on GPGPUs	2702
<i>Nan Jiang, Hengshan Yue, Jingweijia Tan, Mengting Zhou, Xiaonan Wang, Yuchun Wang, Wenda Wei, Meikang Qiu, Xiaohui Wei</i>	
Construction of DAG Models for Autonomous Systems	2709
<i>Jing Huang, Kuan Jiang, Weijie Wang, Wei Liang, Wanli Chang</i>	
ReaLM: Reliable and Efficient Large Language Model Inference with Statistical Algorithm-Based Fault Tolerance	2715
<i>Tong Xie, Jiawang Zhao, Zishen Wan, Zuodong Zhang, Yuan Wang, Runsheng Wang, Ru Huang, Meng Li</i>	
InterConFuzz: A Fuzzing-Based Comprehensive NoC Verification Framework	2722
<i>Samit Shahnawaz Miftah, Hyunmin Kim, Kanad Basu</i>	
Multicore Environment State Representation for Agent-Directed Test Generation.....	2729
<i>Bruno D. Miranda, Luiz M. V. Pereira, Márcio Castro, Luiz C. V. Dos Santos</i>	

GSIM: Accelerating RTL Simulation for Large-Scale Designs.....	2736
<i>Lu Chen, Dingyi Zhao, Zihao Yu, Ninghui Sun, Yungang Bao</i>	
GEM: GPU-Accelerated Emulator-Inspired RTL Simulation	2743
<i>Zizheng Guo, Yanqing Zhang, Runsheng Wang, Yibo Lin, Haoxing Ren</i>	
Simulation-Based Parallel Sweeping: A New Perspective on Combinational Equivalence Checking	2750
<i>Tianji Liu, Evangeline F. Y. Young</i>	
Parallel Dynamic Partitioning for Datapath Combinational Equivalence Checking	2757
<i>Shuai Zhou, Weikang Zhang, Xindi Zhang, Zite Jiang, Haihang You, Shaowei Cai</i>	
Blaze: An Efficient Bit-Sparse Attention Architecture with Workload Orchestration Optimization	2764
<i>Runzhou Zhang, Faxian Sun, Yiming Wang, Kunchen Zou, Zhinan Qin, Jianli Chen, Jun Yu, Kun Wang</i>	
DenSparSA: A Balanced Systolic Array Approach for Dense and Sparse Matrix Multiplication	2771
<i>Ziheng Wang, Ruiqi Sun, Xin He, Tianrui Ma, An Zou</i>	
UniCoS: A Unified Neural and Accelerator Co-Search Framework for CNNs and ViTs	2778
<i>Wei Fu, Wenqi Lou, Cheng Tang, Hongbing Wen, Yunji Qin, Lei Gong, Chao Wang, Xuehai Zhou</i>	
GSAcc: Accelerate 3D Gaussian Splatting Via Depth Speculation and Gaussian-Centric Rasterization	2784
<i>Mengtian Yang, Yipeng Wang, Chieh-Pu Lo, Xiuhao Zhang, Sirish Oruganti, Jaydeep P. Kulkarni</i>	
High-Throughput Point-Cloud Accelerator with Sparsity-Aware Hierarchical Neighbor Voxel Search and Skipping	2791
<i>Yun-Chia Yu, Suraj Pn Reddy, Aryan Devrani, Anirudh Srinivasan, Saianudeep Reddy Nayini, Sohyeon Kim, Sung-Joon Jang, Sang-Seol Lee, Mingu Kang</i>	
April: Accuracy-Improved Floating-Point Approximation for Neural Network Accelerators.....	2798
<i>Yonghao Chen, Jiaxiang Zou, Xinyu Chen</i>	
CVMAX: Accelerator Architecture with Polar Form Multiplication for Complex-Valued Neural Networks	2805
<i>Hyunwuk Lee, Sungbin Kim, Sungwoo Kim, Won Woo Ro</i>	
An Efficient Bit-Level Sparse MAC-Accelerated Architecture with SW/HW Co-Design on FPGA.....	2812
<i>Chenming Zhang, Lei Gong, Chao Wang, Xuehai Zhou</i>	
IG-Kway: Incremental K-Way Graph Partitioning on GPU	2819
<i>Wan Luan Lee, Shui Jiang, Dian-Lun Lin, Che Chang, Boyang Zhang, Yi-Hua Chung, Ulf Schlichtmann, Tsung-Yi Ho, Tsung-Wei Huang</i>	
FireGuard: A Generalized Microarchitecture for Fine-Grained Security Analysis on OoO Superscalar Cores	2826
<i>Zhe Jiang, Sam Ainsworth, Timothy Jones</i>	
EdgeMM: Multi-Core CPU with Heterogeneous AI-Extension and Activation-Aware Weight Pruning for Multimodal LLMs at Edge	2833
<i>Kangbo Bai, Le Ye, Ru Huang, Tianyu Jia</i>	

ACRS: Adjacent Computation Resource Sharing Among Partitioned GPU Sub-Cores.....	2840
<i>Penghao Song, Chongxi Wang, Chenji Han, Haoyu Zhao, Tingting Zhang, Tianyi Liu, Jian Wang</i>	
Swift Or Exact? Boosting Efficient Microarchitecture DSE Via Multi-Fidelity Partial Order Prediction	2847
<i>Hang Liu, Hao Geng, Zhuolun He, Qi Sun, Cheng Zhuo</i>	
Principle-Based Dataflow Optimization for Communication Lower Bound in Operator-Fused Tensor Accelerator.....	2854
<i>Lei Xu, Chen Yin, Zelong Yuan, Weiguang Sheng, Jianfei Jiang, Qin Wang, Naifeng Jing</i>	
GoPTX: Fine-Grained GPU Kernel Fusion by PTX-Level Instruction Flow Weaving.....	2861
<i>Kan Wu, Zejia Lin, Mengyue Xi, Zhongchun Zheng, Wenxuan Pan, Xianwei Zhang, Yutong Lu</i>	
DARIS: An Oversubscribed Spatio-Temporal Scheduler for Real-Time DNN Inference on GPUs	2868
<i>Amir Fakhim Babaei, Thidapat Chantem</i>	
Late Breaking Results: A Diffusion-Based Framework for Configurable and Realistic Multi-Storage Trace Generation	2875
<i>Seohyun Kim, Junyoung Lee, Jongho Park, Jinhyung Koo, Sungjin Lee, Yeseong Kim</i>	
Late Breaking Results: A Fast Nearest Neighbor Search Acceleration for 3D Point Cloud.....	2877
<i>Jiniao Li, Teng Wang, Qianyu Cheng, Zhendong Zheng, Lei Gong, Chao Wang, Xi Li, Xuehai Zhou</i>	
Late Breaking Results: A Geometric Diffusion Model for Macro Placement Generation.....	2879
<i>Jongho Yoon, Jinsung Jeon, Seokhyeong Kang</i>	
Late Breaking Results: Advanced PCB Placement with Irregular Components for Efficient Collision Detection and Routability Optimization	2881
<i>Chien-Hao Tsou, Zhu-Xun Lee, Yao-Wen Chang</i>	
Late Breaking Results: An Efficient and Scalable Track Assignment with GPU Parallelism	2883
<i>Genggeng Liu, Pengcheng Huang, Zepeng Li, Wen-Hao Liu, Xing Huang, Wenzhong Guo</i>	
Late Breaking Results: Automated Topology Generation for Power Amplifier Designs Through BiLSTM-Based DNN and Multi-Objective Optimizations	2885
<i>Lida Kouhalvandi, Sercan Aygun, M. Hassan Najafi, Arman Roohi</i>	
Late Breaking Results: BLAST: Bisection-Free Learning Approach for Statistical Timing Characterization.....	2887
<i>Kai Jing, Tao Bai, Zeyuan Deng, Junming Jiao, Peng Cao</i>	
Late Breaking Results: Breaking Symmetry— Unconventional Placement of Analog Circuits Using Multi-Level Multi-Agent Reinforcement Learning.....	2889
<i>Supriyo Maji, Linran Zhao, Souradip Poddar, David Z. Pan</i>	
Late Breaking Results: Customized Diffusion Model Empowered by Heterogeneous Graph Network for Effective Floorplanning	2891
<i>Xinglin Zheng, Hao Gu, Keyu Peng, Youwen Wang, Wenxing Zhu, Ziran Zhu</i>	
Late Breaking Results: Decentralized Voting-Based Attestation for IoT Devices.....	2893
<i>Mohamed Alsharkawy, Eren Sönmez, Jeferson González-Gómez, Hassan Nassar, Jörg Henkel</i>	

Late Breaking Results: Encoder-Decoder Generative Diffusion Transformer Towards Push-Button Analog IC Sizing	2895
<i>Filipe Azevedo, Nuno Lourenço, Ricardo Martins</i>	
Late Breaking Results: Fine-Tuning LLMs for Test Stimuli Generation	2897
<i>Hyeonwoo Park, Seonghyeon Park, Seokhyeong Kang</i>	
Late Breaking Results: FPGen-3D: Automated Framework for 3D-FPGA Architecture Generation and Exploration	2899
<i>Ismael Youssef, Cong Callie Hao</i>	
In-Memory Arithmetic: Enabling Division with Stochastic Logic	2901
<i>Farzad Razi, Mehran Shoushtari Moghadam, M. Hassan Najafi, Sercan Aygun, Marc Riedel</i>	
Late Breaking Results: Hybrid Logic Optimization with Predictive Self-Supervision	2903
<i>Rongliang Fu, Ran Zhang, Ziyang Zheng, Zhengyuan Shi, Yuan Pu, Junying Huang, Qiang Xu, Tsung-Yi Ho</i>	
Late Breaking Results: Less Sense Makes More Sense: In-Sensor Compressive Learning for Efficient Machine Vision	2905
<i>Yiwen Liang, Weidong Cao</i>	
Late Breaking Results: Multi-Objective Multi-Bit Flip-Flop Placement Considering Pre-Placed Cells	2907
<i>Cheng-Yen Li, Chuan-Chi Su, Zheng-Wei Chen, Shao-Hsiang Chen, Yao-Wen Chang</i>	
Late Breaking Results: Novel Design of MTJ-Based Unified LIF Spiking Neuron and PUF	2909
<i>Milad Tanavardi Nasab, Wu Yang, Himanshu Thapliyal</i>	
Late Breaking Results: On-The-Fly Hadamard Hypervector Processing for Efficient Hyperdimensional Computing	2911
<i>Abu Kaisar Mohammad Masum, Mehran Shoushtari Moghadam, Sabrina Hassan Moon, Ahmed Mamdouh Mohamed Ahmed, M. Hassan Najafi, Dayane Reis, Sercan Aygun</i>	
Late Breaking Results: Opera: An Open and Efficient Platform for Data-Driven Synthesis of Analog Circuits	2913
<i>Shikai Wang, Yaolong Hu, Zhiqiang Yi, Taiyun Chi, Weidong Cao</i>	
Late Breaking Results: Scalable GPU-Friendly Parallelization for Sweep-Based Maze Routing	2915
<i>Cheng-Yu Chiang, Zong-Ying Cai, Chao-Chi Lan, Yan-Jen Chen, Yang Hsu, Yao-Wen Chang, Hung-Ming Chen</i>	
Late Breaking Results: Source-Aware Adaptive Cache Management for CXL-Enabled Disaggregated Memory Sharing	2917
<i>Qianyu Cheng, Jiajun Ji, Teng Wang, Zihan Wang, Lei Gong, Chao Wang, Xuehai Zhou</i>	
Late Breaking Results: Statistical Timing Graph Scheduling Algorithm for GPU Computation	2919
<i>Chih-Chun Chang, Tsung-Wei Huang</i>	
Late Breaking Results: The Hidden Risks of Activation Duration in PLPUFs	2921
<i>Mohamed Alsharkawy, Jan Zwerschke, Hassan Nassar, Jeferson González-Gómez, Jörg Henkel</i>	
Late Breaking Results: Utilization of Hybrid Threshold-Voltage Flip-Flops for Power Recovery	2923
<i>Sehyeon Chung, Hyun-Chul Hwang, Byung Su Kim, Jaeha Lee, Kunhyuk Kang, Taewhan Kim</i>	

Late Breaking Results: Versatile 4:1 Multiplexer Using 1T1R RRAM Crossbar for High Speed In-Memory Computing	2925
<i>B Vinodh Kumar, Binsu J Kailath</i>	
Late Breaking Results: Warpage-Aware Generative Floorplanning for Reliable Advanced Packaging	2927
<i>Min-Hung Chen, Cheng-Yen Li, Chuan-Chi Su, Yao-Wen Chang, Tung-Chieh Chen</i>	
LLMs: A Driving Force in Next Generation Digital Design Automation.....	2929
<i>Matheus T. Moreira, Aram H. Markosyan, Chris Cummins, Warren Hunt, Gabriel Synnaeve, Edith Beigne</i>	
LLMs Meet Post-Silicon Test Engineering: A New Era (Invited)	2935
<i>Li-C. Wang</i>	
Invited Paper: Enhancing Design Automation with Quantum Algorithms for Chip Design	2940
<i>Gabriel P. L. M. Fernandes, Matheus S. Fonseca, Amanda G. Valério, Nicolás A. C. Carpio, João Marcus Epifânio Morais De Assunção, Rafael V. Aroca, Celso J. Villas-Boas, Dario Sassi Thober</i>	
Invited: SambaNova SN40L: Unleashing Agentic AI with Dataflow	2945
<i>Raghu Prabhakar, Pushkar Nandkar, Darshan Gandhi, Nasim Farahini, Hakan Zeffer</i>	
Invited: Enhancing Test Quality by Targeting Timing Marginalities Due to Process Variations.....	2950
<i>Adit D. Singh, Mukarram Ali Faridi</i>	
Device-Aware Test: A Means to Attack Unmodelled Defects (Invited)	2954
<i>Said Hamdioui, Mottaqiallah Taouil</i>	
Invited: EDA for Heterogeneous Integration.....	2958
<i>Emad Haque, Pragnya Nalla, Chetan Choppali Sudarshan, Divya Yogi, Hangyu Zhang, Chaitali Chakrabarti, Vidya A. Chhabria, Ramesh Harjani, Jeff Zhang, Sachin S. Sapatnekar</i>	
Towards Secure Data Management Using Multi-Cryptographic Solutions (Invited).....	2962
<i>Shufan Zhang, Xi He, Ashish Kundu, Sujaya Maiyya, Sharad Mehrotra, Shantanu Sharma</i>	
Security Opportunities and Challenges for Disaggregated Architectures (Invited).....	2966
<i>Elisa Bertino, Imtiaz Karim, Ashish Kundu</i>	
Quantum-Resistant Security: PQC Readiness and Research Challenges (Invited)	2970
<i>Ashish Kundu, Ramana Kompella</i>	

Author Index