

2026 31st Asia and South Pacific Design Automation Conference (ASP-DAC 2026)

**Lantau, Hong Kong
19-22 January 2026**

Pages 1-729



**IEEE Catalog Number: CFP26ASP-POD
ISBN: 979-8-3315-9124-3**

**Copyright © 2026 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:	CFP26ASP-POD
ISBN (Print-On-Demand):	979-8-3315-9124-3
ISBN (Online):	979-8-3315-9123-6
ISSN:	2153-6961

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

Video-based Visible-Event Cross-modal Person Re-identification for Edge AI Surveillance Systems.....	1
<i>Xinyun Zhang, Zixiao Wang, Yurui Kuang, Bei Yu</i>	
HeteroSTA: A CPU-GPU Heterogeneous Static Timing Analysis Engine with Holistic Industrial Design Support.....	8
<i>Zizheng Guo, Haichuan Liu, Xizhe Shi, Shenglu Hua, Zuodong Zhang, Chunyuan Zhao, Runsheng Wang, Yibo Lin</i>	
REDM: Regression-Guided Diffusion Modeling for Universal Soft Sensor Enhancement in Semiconductor Process Control.....	15
<i>Weiping Xie, Yumeng Shi, Pang Guo, Yining Chen</i>	
Benchmarking Continual Learning on Netlists with Circuit-Targeted Graph Neural Networks	22
<i>Rupesh Raj Karn, Johann Knechtel, Ozgur Sinanoglu</i>	
LiveHPS-Lite: A Lightweight LiDAR-based Motion Capture System for Edge Applications.....	29
<i>Yiren Zhu, Junsheng Zhou, Yiming Ren, Hanshu Hezi, Yuexin Ma, Xin Lou</i>	
ML-driven Design Technology Co-Optimization Framework for Advanced Technology Nodes	36
<i>Hyunbae Seo, Handong Cho, Sehyeon Chung, Kyumyung Choi, Taewhan Kim</i>	
Standard Cell Layout Generation: Methodological Evolution and Architectural Impacts	42
<i>Junghyun Yoon, Ikkyum Kim, Gyumin Kim, Sojung Park, Heechun Park</i>	
Fast Timing Library Characterization Through Selective Use of Regression Models.....	49
<i>Manikanta Prahlad Manda, Seunggyu Lee, Daijoon Hyun</i>	
HyFault: Targeted Fault Injection Attacks on Hyperdimensional Computing Accelerators	56
<i>Brojogopal Sapui, Mehdi B. Tahoori</i>	
PIR-Cache: Mitigating Conflict-Based Cache Side-Channel Attacks via Partial Indirect Replacement	63
<i>Hao Ma, Zhidong Wang, Wei Song</i>	
An Efficient Defense Method Based on Progressive Fault-Aware Training and JS Divergence-Guided TMR for DNNs against Bit-Flip Attacks	70
<i>Huarun Zhou, Ran Dong, Zhaohui Guo, Qiang Liu</i>	
X-Matrix Shield: Defeating Tilted FIB and Rerouting Attacks through 3D-Interlaced Protection	77
<i>Jiaji He, Junfeng Cai, Yaohua Wang, Fei Zhao, Mao Ye, Yongqiang Lyu</i>	
Hardware-Efficient Union-Find Decoder Towards Scalable Topological Quantum Codes.....	84
<i>Shuang Liang, Jubo Xu, Yuncheng Lu, Hao Mark Chen, Bo Yuan, Hongxiang Fan</i>	
Reinforcement Learning for Enhanced Advanced QEC Architecture Decoding.....	90
<i>Yidong Zhou, Lingyi Kong, Yifeng Peng, Zhiding Liang</i>	
Quantum Instruction Set Architecture: The Good, the Bad, and the Future	97
<i>Jianxin Chen</i>	
Old School Never Die: A Classic Yet Novel Algorithm for Computing RC Current Response in VLSI.....	98
<i>Zongfeng Ma, Zhong Guan</i>	

TargetFuzz: Enabling Directed Graybox Fuzzing via SAT-Guided Seed Generation.....	105
<i>Raghul Saravanan, Sai Manoj Pudukotai Dinakarrao</i>	
VeriRAG: A Knowledge Graph-Augmented RAG for Verilog and Assertion Generation	112
<i>Jayanth Thangellamudi, Raghul Saravanan, Sai Manoj Pudukotai Dinakarrao</i>	
Spiking-NeRF: Neural Graphics Acceleration With Spiking Feature Encoding for Edge 3D Rendering	119
<i>Jianzhen Gao, Wei Liu, Yue Liu, Hengyi Zhou, Zhiyi Yu, Shanlin Xiao</i>	
FlowQ: Fixed-point Low-precision Post-Training Quantization Framework for Efficient and Accurate SNN Inference.....	126
<i>Faaiz Asim, Sanhtet Aung, Jongeun Lee</i>	
An Algorithm-Hardware Co-Design for Efficient and Robust Spiking Neural Networks via Sparsity	133
<i>Wei Liu, Yinsheng Chen, Jilong Luo, Yusa Wang, Zhiyi Yu, Shanlin Xiao</i>	
LOKI: a 0.266 pJ/SOP Digital SNN Accelerator with Multi-Cycle Clock-Gated SRAM in 22 nm.....	140
<i>Rick Luiken, Lorenzo Pes, Manil Dev Gomony, Sander Stuijk</i>	
PipeViT: Accelerating Vision Transformers via Intra-Layer Pipelining	147
<i>Xilang Zhou, Yiheng Xu, Haodong Lu, Jun Yu, Kun Wang</i>	
ConfASR: A Conformer Block Accelerator for Speech Recognition Optimized for Edge Devices.....	154
<i>Malte Wabnitz, Max Nilovic, Finn Scholz, Dominik Friedrich, Christian Lanius, Jie Lou, Tobias Gemmeke</i>	
LAPA: Log-Domain Prediction-Driven Dynamic Sparsity Accelerator for Transformer Model	161
<i>Huizheng Wang, Hongbin Wang, Shaojun Wei, Yang Hu, Shouyi Yin</i>	
BitStopper: An Efficient Transformer Attention Accelerator via Stage-fusion and Early Termination	169
<i>Huizheng Wang, Hongbin Wang, Shaojun Wei, Yang Hu, Shouyi Yin</i>	
JaneEye: A 12-nm 2K-FPS 18.9- μ J/Frame Event-based Eye Tracking Accelerator.....	177
<i>Tao Han, Ang Li, Qinyu Chen, Chang Gao</i>	
Scalable Optimization with GIS-PIM: A Generalized Integer-State Probabilistic Ising Machine.....	184
<i>Chirag Garg, Sayeef Salahuddin</i>	
A Scalable and High-Quality Qubit Mapping and Shuttling Framework for Neutral Atom Quantum Devices	191
<i>Sung-Ying Hsieh, Wai-Kei Mak</i>	
Quantum Oracle Synthesis from HDL Designs via Multi Level Intermediate Representation	198
<i>Giacomo Lancellotti, Filippo Buda, Giacomo Carugati, Daniele Gazzola, Alessandro Barenghi, Giovanni Agosta, Gerardo Pelosi</i>	
Survival of the Optimized: An Evolutionary Approach to T-depth Reduction	205
<i>Archisman Ghosh, Avimta Chatterjee, Swaroop Ghosh</i>	
Subgraph-based Qubit Mapping for Noisy Intermediate-Scale Quantum Computing	212
<i>Junyi Gao, Wei-Hsiang Tseng, Yao-Wen Chang</i>	
A 100V 86.2% Efficiency Fibonacci-Dickson Hybrid Boost Converter for Acoustic Screen Applications.....	220
<i>Chen Hu, Yifan Jiang, Yan Lu, Junmin Jiang</i>	

A 5-to-1V DLDO-Hybrid-Sigma Converter Achieving Fast Transient for High-Density Power Delivery	223
<i>Zizhe Huang, Yuxiang Li, Yuekang Guo, Jing Jin, Jianjun Zhou, Junmin Jiang</i>	
Full-Stack System Design and Prototyping for Fully Programmable Electronic-Photonic Neurocomputing	226
<i>Yinyi Liu, Peiyu Chen, Bohan Hu, Wei Zhang, Jiang Xu</i>	
Analysis and Design of Oblong Coils and Standard-Cell-Based Receiver for Area-Efficient Edge-Coupled Inductive Coupling Transceiver	230
<i>Yuki Mitarai, Mototsugu Hamada, Atsutake Kosuge</i>	
A RHP-Zero-Free Hybrid Step-Up Converter With 95.1% Peak Efficiency for Fast-Transient Applications.....	234
<i>Junyi Ruan, Junmin Jiang, Chenzhou Ding, Yanhui Wu, Ka Nang Leung, Xun Liu</i>	
A Relaxation Oscillator with 2.93 μ J/cycle Energy Efficiency and 0.068% Period Jitter	237
<i>Yongjuan Shi, Xun Liu, Chen Hu, Xiyuan Tang, Junmin Jiang</i>	
TFLOP: Towards Energy-Efficient LLM Inference An FPGA-Affinity Accelerator with Unified LUT-based Optimization	240
<i>Zongwu Wang, Zhongyi Tang, Fangxin Liu, Chenyang Guan, Li Jiang, Haibing Guan</i>	
MF-ECC: Memory-Free Error Correction for Hyperdimensional Computing Edge Accelerators	244
<i>Mahboobe Sadeghipour Roodsari, Mahta Mayahinia, Mehdi B. Tahoori</i>	
Thermo-NAS: Thermal-resilient ultralow-cost IGZO-based Flexible Neuromorphic Circuits	251
<i>Priyanjana Pal, Tara Gheshlaghi, Suman Balaji, Emre Ozer, Mehdi B. Tahoori</i>	
Flawase: A SET Fault Injection Framework Towards Exhaustive System-Level Impact Evaluation	258
<i>Mingtao Zhang, Quan Cheng, Masanori Hashimoto</i>	
MASS: A Masking-aware Search Framework for Reliable QC-LDPC Code Construction in SSDs	265
<i>Xiaolu Li, Dingxin Wang, Zhengyao Ding, Jinye Wu, Qingnan Hu, Patrick P. C. Lee, Yuchong Hu, Dan Feng</i>	
TIMBER: A Fast Algorithm for Timing and Power Optimization using Multi-bit Flip-flops	273
<i>Aditya Das Sarma, Shui Jiang, Wan Luan Lee, Tsung-Yi Ho, Tsung-Wei Huang</i>	
GENIAL: Generative Design Space Exploration via Network Inversion for Low Power Algorithmic Logic Units.....	280
<i>Maxence Bouvier, Ryan Amaudruz, Felix Arnold, Renzo Andri, Lukas Cavigelli</i>	
REvolution: An Evolutionary Framework for RTL Generation driven by Large Language Models.....	289
<i>Kyungjun Min, Kyumin Cho, Junhwan Jang, Seokhyeong Kang</i>	
AC-Refiner: Efficient Arithmetic Circuit Optimization Using Conditional Diffusion Models	296
<i>Chenhao Xue, Kezhi Li, Jiaxing Zhang, Yi Ren, Zhengyuan Shi, Chen Zhang, Yibo Lin, Lining Zhang, Qiang Xu, Guangyu Sun</i>	
DeepCut: Structure-Aware GNN Framework for Efficient Cut Timing Prediction in Logic Synthesis.....	304
<i>Lingfeng Zhou, Yilong Zhou, Hao Gong, Zhengyuan Shi, Qiang Xu, Zhufei Chu, Yue Wu, Xiaoyan Yang</i>	
Lorecast: Layout-Aware Performance and Power Forecasting from Natural Language	311
<i>Runzhi Wang, Prianka Sengupta, Cristhian Roman-Vicharra, Yiran Chen, Jiang Hu</i>	

DeepPiC: xPU-PIM Cluster Architecture with Adaptive Resource-Aware Task Orchestration for DeepSeek-Style MoE Inference	318
<i>Zixu Li, Manni Li, Zijian Huang, Jiayu Yang, Wending Zhao, Yinyin Lin, Chengchen Wang, Haidong Tian, Xiankui Xiong</i>	
MoEA: A Mixture of Experts Accelerator with Direct Token Access and Dynamic Expert Scheduling	325
<i>Zifeng Zhao, Jiewen Zheng, Tianxing Xie, Xinghao Zhu, Gengsheng Chen</i>	
MoEA: A Mixed-Precision Edge Accelerator for CNN-MSA Models with Fine-Tuning Support.....	332
<i>Qiwei Dang, Chengyu Ma, Zhiwang Huo, Guoming Yang, Tian Xia, Wenzhe Zhao, Pengju Ren</i>	
Efficient CPU-GPU Collaborative Inference for MoE-based LLMs on Memory-Limited Systems.....	340
<i>En-Ming Huang, Li-Shang Lin, Chun-Yi Lee</i>	
BOA-3DGS: Backward-Striding Optimized Accelerator for Reduced Memory Contention in 3D Gaussian Splatting Training	348
<i>Hyukjun Kweon, Jongyeop Kim, Jeongwoo Park</i>	
Scalarium: A Unified Scala-based Co-Simulation Framework for Agile Chip Development	355
<i>Yuefeng Zhang, Cheng Zhang, Wenkai Zhou, Binzhe Yuan, Junsheng Chen, Xiangyu Zhang, Hao Geng, Xin Lou</i>	
VFlow: Discovering Optimal Agentic Workflows for Verilog Generation.....	362
<i>Yangbo Wei, Zhen Huang, Lei He, Li Huang, Ting-Jung Lin, Wei W. Xing</i>	
SemanticBBV: A Semantic Signature for Cross-Program Knowledge Reuse in Microarchitecture Simulation	369
<i>Zhenguo Liu, Chengao Shi, Chen Ding, Jiang Xu</i>	
CausalTuner: Will Causality Help High-Dimensional EDA Tool Parameter Tuning	377
<i>Ziyang Yu, Peng Xu, Su Zheng, Siyuan Xu, Hao Geng, Bei Yu, Martin Wong</i>	
Synergistic Bayesian Optimization and Reinforcement Learning with Bidirectional Interaction for Efficient VLSI Constraint Tuning	384
<i>Jiayi Tu, Jindong Tu, Yuxuan Cai, Yi Zhang, Meng Zhang, Tinghuan Chen</i>	
Mastering the Exponential Complexity of Exact Physical Simulation of Silicon Dangling Bonds	391
<i>Willem Lambooy, Jan Drewniok, Marcel Walter, Robert Wille</i>	
Built-In Self-Test for Locating Leakage Defects on Continuous-Flow Microfluidic Chips.....	398
<i>Jiahui Peng, Mengchu Li, Tsun-Ming Tseng, Ulf Schlichtmann</i>	
Accessible Ratio-Specific Mixing: Single-Pressure-Driven Multi-Reagent Mixer Design and Synthesis for 3D-Printed Microfluidics.....	405
<i>Yushen Zhang, Debraj Kundu, Tsun-Ming Tseng, Sudip Roy, Shigeru Yamashita, Ulf Schlichtmann</i>	
ENLighten: Lighten the Transformer, Enable Efficient Optical Acceleration.....	412
<i>Hanqing Zhu, Zhican Zhou, Shupeng Ning, Xuhao Wu, Ray T. Chen, Yating Wan, David Z. Pan</i>	
DCPPC: Digital Computation in Programmable Photonic Circuits	420
<i>Jun-Wei Liang, Iris Hui-Ru Jiang, Kai-Hsiang Chiu</i>	

SuperSAGA: A Supervisor-Subordinate Agentic workflow for the Generation of Assertions	427
<i>Subhajit Paul, Ansuman Banerjee, Sumana Ghosh, Sudhakar Surendran, Raj Kumar Gajavelly</i>	
Understanding and Predicting Vmin Failures in Power Delivery Networks through Multi-Order Droop Signatures	433
<i>Songyu Sun, Jingchao Hu, Zhou Jin, Cheng Zhuo</i>	
AssertMiner: Module-Level Spec Generation and Assertion Mining using Static Analysis Guided LLMs	439
<i>Hongqin Lyu, Yonghao Wang, Jiaxin Zhou, Zhiteng Chao, Tiancheng Wang, Huawei Li</i>	
LLM-Assisted Circuit Verification: A Comprehensive Survey	446
<i>Hongduo Liu, Yuntao Lu, Mingjun Wang, Xufeng Yao, Bei Yu</i>	
HeteroLatch: A CPU-GPU Heterogeneous Latch-Aware Timing Analysis Engine	454
<i>Xizhe Shi, Zizheng Guo, Yibo Lin, Zuodong Zhang, Yun Liang, Runsheng Wang</i>	
Novel Multi-Corner Delay Padding using Path Relationship Analysis and Dual Decomposition	461
<i>Kaixiang Zhu, Jiangnan Li, Lingli Wang, Wai-Shing Luk</i>	
GNN-Based Timing Yield Prediction From Statistical Static Timing Analysis	468
<i>Chenbo Xi, Liang Zhang, Biwei Xie, Pingqiang Zhou</i>	
MIMIC: Machine Intelligence for Scalable Generation of Synthetic Timing Cone Datasets	475
<i>Ajay Yadav, Vinodh Kumar Ramasamy, Juan Arturo Garza, Taylor Hannan, Mark Lee, Mahesh Sharma, Vidya A. Chhabria</i>	
Differentiable Tier Assignment for Timing and Congestion-Aware Routing in 3D ICs	482
<i>Yuan-Hsiang Lu, Hao-Hsiang Hsiao, Yi-Chen Lu, Haoxing Ren, Sung Kyu Lim</i>	
Automatic Recursion Elimination using Recurrence Relations for Synthesis of Stack-free Hardware	489
<i>Adam Musa, Christophe Dubach</i>	
FIFOAdvisor: A DSE Framework for Automated FIFO Sizing of High-Level Synthesis Designs	496
<i>Stefan Abi-Karam, Rishov Sarkar, Suhail Basalama, Jason Cong, Callie Hao</i>	
HLS-Timer: Fine-Grained Path-Level Timing Estimation for High-Level Synthesis	503
<i>Zibo Hu, Zhe Lin, Renjing Hou, Xingyu Qin, Jianwang Zhai, Kang Zhao</i>	
FESTAL: Dataflow Accelerator Synthesis Framework with Graph-Based Fusion for FPGA	510
<i>Ruifan Xu, Yuyang Zou, Yun Liang</i>	
ARCS: Architecture-Responsive CGRA Scheduling	519
<i>Omar Ragheb, Jason Anderson</i>	
CoLoRA: A Collaborative Scheduling Framework for Multi-Tenant LoRA LLM Inference	526
<i>Zechao Lin, Xingbin Wang, Yiming Xie, Dan Meng, Rui Hou</i>	
AutoVeriFix: Automatically Correcting Errors and Enhancing Functional Correctness in LLM-Generated Verilog Code	533
<i>Yan Tan, Xiangchen Meng, Zijun Jiang, Yangdi Lyu</i>	
HDLxGraph: Bridging Large Language Models and HDL Repositories via HDL Graph Databases	540
<i>Pingqing Zheng, Jiayin Qin, Fuqi Zhang, Niraj Chitla, Zishen Wan, Shang Wu, Yu Kevin Cao, Caiwen Ding, Yang Katie Zhao</i>	

Chat-A2: An LLM-aided Design Space Exploration Framework for High-Performance CPU Design	547
<i>Zhantong Zhu, Kangbo Bai, Tianyu Jia</i>	
BitROM: Weight Reload-Free CiROM Architecture Towards Billion-Parameter 1.58-bit LLM Inference	554
<i>Wenlun Zhang, Xinyu Li, Shimpei Ando, Kentaro Yoshioka</i>	
BALANCE: Bit and Layer-Aware Lightweight ECC Design Method for In-Flash Computing Based LLM Inference Accelerator	561
<i>Changwei Yan, Lishuo Deng, Mingbo Hao, Cai Li, Weiwei Shan</i>	
BLADE: Boosting LLM Decoding's Communication Efficiency in DRAM-based PIM	568
<i>Yilong Zhao, Fangxin Liu, Zongwu Wang, Mingjian Li, Mingxing Zhang, Chixiao Chen, Li Jiang</i>	
SpAct-NDP: Efficient LLM Inference via Sparse Activation on NDP-GPU Heterogeneous Architecture	575
<i>Jiaming Xu, Tongxin Xie, Yongkang Zhou, Jinhao Li, Yaoxiu Lian, Zhenhua Zhu, Yu Wang, Guohao Dai</i>	
Quantifying Compiler-induced Reliability Loss in Software-Implemented Hardware Fault Tolerance	582
<i>Davide Baroffio, Johannes Geier, Federico Reghenzani, Ulf Schlichtmann, William Fornaciari</i>	
WARP: Workload-Aware Reference Prediction for Reliable Multi-Bit FeFET Readout under Charge-Trapping Degradation*	590
<i>Dhruv Thapar, Ashish Reddy Bommana, Arjun Chaudhuri, Kai Ni, Krishnendu Chakrabarty</i>	
PV-ReCAM: Process Variation-Aware Testing for ReRAM-based Content Addressable Memory.....	597
<i>Haneen G. Hezayyin, Mahta Mayahinia, Mehdi Tahoori</i>	
When Posit Meets Microscaling: Energy Efficient Posit-Based Processing Element for Edge AI Computation	604
<i>Yulin Wang, Seok-Bum Ko, Qi Wen, Zhiqiang Wei, Hao Zhang</i>	
When Low-Rank Meets Mixed-Precision: Training-Free Joint Compression for Efficient LLM Inference	611
<i>Junjie Wang, Fangxin Liu, Jinqi Zhu, Chenyang Guan, Tao Yang, Li Jiang, Haibing Guan</i>	
Precision-Scalable Microscaling Datapaths with Optimized Reduction Tree for Efficient NPU Integration	618
<i>Stef Cuyckens, Xiaoling Yi, Robin Geens, Joren Dumoulin, Martin Wiesner, Chao Fang, Marian Verhelst</i>	
Understand and Detect: Lithographic Hotspot Detection by the Interpretable Graph Attention Network	625
<i>Andy Liu, Silin Chen, Guohao Wang, Wenzheng Zhao, Yuxiang Fu, Ningmu Zou</i>	
Beyond Labels: Data-Efficient Wafer Yield Prediction with TabESA.....	632
<i>Pang Guo, Yining Chen</i>	
Code, Not Canvas: Multi-Agent Layout Generation Beyond Vision Models.....	639
<i>Haoyu Yang, Haoxing Ren</i>	
Integrated Re-Fragmentation and Curve Correction for Curvilinear Optical Proximity Correction	646
<i>Seohyun Kim, Shilong Zhang, Junha Jang, Youngsoo Shin</i>	

MuaLLM: A Multimodal Large Language Model Agent for Circuit Design Assistance with Hybrid Contextual Retrieval-Augmented Generation	653
<i>Pravallika Abbineni, Saoud Aldowaish, Colin Liechty, Soroosh Noorzad, Ali Ghazizadeh Ghalati, Morteza Fayazi</i>	
MOSTAR: Multi-Stage Hierarchical Bayesian Optimization for Substructure-Aware High-Dimensional Analog Circuit Sizing	660
<i>Weijian Fan, Haoyi Zhang, Weibin Lin, Runsheng Wang, Yibo Lin</i>	
Effective RC Reduction via Graph Sparsification for Accurate Post-Simulation of Mixed-Signal ICs	667
<i>Yibin Zhang, Zhiqiang Liu, Shan Shen, Chao Hu, Wenjian Yu</i>	
CADC: Crossbar-Aware Dendritic Convolution for Efficient In-memory Computing	674
<i>Shuai Dong, Junyi Yang, Ye Ke, Hongyang Shang, Arindam Basu</i>	
M3DKV: Monolithic 3D Gain Cell Memory Enabled Efficient KV Cache & Processing	681
<i>Jiaqi Yang, Yanbo Su, Yihan Fu, Jianshi Tang, Bonan Yan</i>	
MemSearch: An Efficient Memristive In-memory Search Engine with Configurable Similarity Measures	688
<i>Yingjie Yu, Houji Zhou, Jiancong Li, Tong Hu, Jia Chen, Yi Li, Xiangshui Miao</i>	
SpANNS: Optimizing Approximate Nearest Neighbor Search for Sparse Vectors Using Near Memory Processing	695
<i>Tianqi Zhang, Flavio Ponzina, Tajana Rosing</i>	
NBCache: An Efficient and Scalable Non-Blocking Cache for Coherent Multi-Chiplet Systems	703
<i>Zhirong Ye, Yongchang Zhang, Peilin Wang, Tao Lu, Zhaolin Li, Zhiyi Yu, Mingyu Wang</i>	
A Low-Power 12-lead Arrhythmia Detection SoC Featuring a Reconfigurable CNN and Mixed-Precision Computing	710
<i>Yuejun Zhang, Hanyu Shi, Qikang Li, Huihong Zhang, Xinyu Li, Qingxin Xie, Zhenkai Zhou, Pengjun Wang</i>	
A Precision-Scalable Accelerator for Compressive Hyperspectral Image Reconstruction with a Lightweight DUN	716
<i>Shuo Zhang, Shengzhi Qiang, Wendong Mao, Zhongfeng Wang</i>	
Enhancing Trustworthiness with Mixed Precision: Benchmarks, Opportunities, and Challenges	723
<i>Guanxi Lu, Hao Mark Chen, Zhiqiang Que, Wayne Luk, Hongxiang Fan</i>	
GEmFuzz: Uncovering System-Level Vulnerabilities in SoCs via Emulation-Based Grey-Box Fuzzing	730
<i>Shuvagata Saha, Ahmed Alhurubi, Tanvir Rahman, Hasan Al Shaikh, Sujan Kumar Saha, Farimah Farahmandi, Mark Tehranipoor</i>	
Pack Defender: Proactive Defense Against Packet Attacks in NoCs Using an XGBoost-RNN Model	737
<i>Shengkai Hu, Haoyu Wang, Basel Halak, Boojoong Kang</i>	
Silentflow: Leveraging Trusted Execution for Resource-Limited MPC via Hardware-Algorithm Co-design	744
<i>Zhuoran Li, Hanieh Totonchi Asl, Ebrahim Nouri, Yifei Cai, Dan Zhao</i>	
ANIMo: Accelerating Nested Isolation with Monitor-free Domain Transition	751
<i>Yibin Xu, Han Wang, Yue Jin, Tianyi Huang, Tianyue Lu, Mingyu Chen</i>	

Reinforced Logic-Based Distributed Routing within Isolated Secure Zones	758
<i>Yogesh Verma, Mattis Hasler, Sebastian Haas, Friedrich Pauls</i>	
MdpPlanner: Mask-Driven Floorplan via Reinforcement Learning-Based Placement Order	765
<i>Yue Wu, Caiyu Chen, Xiaoyan Yang</i>	
C3PO: Commercial-Quality Global Placement via Coherent, Concurrent Timing, Routability, and Wirelength Optimization	772
<i>Yi-Chen Lu, Hao-Hsiang Hsiao, Rongjian Liang, Wen-Hao Liu, Haoxing Ren</i>	
A Timing-Driven Hierarchical Macro Placement Framework for Large-Scale Complex IP Blocks	779
<i>Wei Fu, Lixin Chen, Jinghui Zhou, Ziran Zhu</i>	
Accelerating Electrostatics-based Global Placement with Enhanced FFT Computation	786
<i>Hangyu Zhang, Sachin S. Sapatnekar</i>	
Comprehensive Delay-Aware Net Weighting Framework for Timing-Driven Global Placement.....	793
<i>Lixin Chen, Keyu Peng, Jinghui Zhou, Hao Gu, Wei Fu, Shuting Cai, Ziran Zhu</i>	
DCLOG: Don't Cares-based Logic Optimization using Pre-training Graph Neural Networks.....	800
<i>Rongliang Fu, Libo Shen, Ziyi Wang, Zhengxing Lei, Zixiao Wang, Junying Huang, Bei Yu, Tsung-Yi Ho</i>	
SOFA-H: Post-Synthesis Area Optimization via Functionally Encoded, Net-Driven Subgraph Mining and SAT-Based Hypercell Remapping.....	807
<i>Jimmy Y.-C. Lee, Yen-Ju Su, Jiun-Cheng Tsai, Aaron C.-W. Liang, Charles H.-P. Wen, Hsuan-Ming Huang</i>	
Formalization of Rectification Learning for Economic Design Updates.....	814
<i>Victor N. Kravets, Jie-Hong R. Jiang</i>	
PigMap3: A Physically Aware Incremental Mapping Framework with On-the-fly Post-Layout Critical Path Tracking.....	822
<i>Hongyang Pan, Cunqing Lan, Zhiang Wang, Xuan Zeng, Fan Yang, Keren Zhu</i>	
CombRewriter: Enabling Combinational Logic Simplification in MLIR-Based Hardware Compiler.....	829
<i>Haisheng Zheng, Zhuolun He, Shuo Yin, Yuzhe Ma, Bei Yu</i>	
SCION: A Comprehensive Simulation Framework for Charge-based In-Memory Computing for Rapid Evaluation of Hardware Non-Idealities and DNN	836
<i>Doug Hyun Kim, Akul Malhotra, Sumeet Kumar Gupta</i>	
SABCIM: Self-Adaptive Biasing Scheme for Accurate and Efficient Analog Compute-in-Memory	844
<i>Yashvardhan Biyani, Abhairaj Singh, Rajendra Bishnoi, Said Hamdioui</i>	
CDACiM: A Charge-Domain Compute-in-Memory Macro for FP/INT MAC Operations with Reconfigurable Capacitor Digital-Analog-Converter.....	851
<i>Jinting Yao, Zeyu Yang, Yuxiao Jiang, Yuxiao Yang, Zheyu Yan, Cheng Zhuo, Xunzhao Yin</i>	
Learnable Center-Based Quantization for Efficient Analog PIM with Reduced ADC Precision	858
<i>Sang Heum Yeon, Jong Hwan Ko</i>	
RL-Guided Thermal-Aware Quantization for Efficient and Robust ReRAM CIM Systems.....	865
<i>Lihua An, Jiayi Li, Pingqiang Zhou</i>	

SnipSnap: A Joint Compression Format and Dataflow Co-Optimization Framework for Efficient Sparse LLM Accelerator Design	872
<i>Junyi Wu, Chao Fang, Zhongfeng Wang</i>	
ALMA: Adaptive Co-optimization of Loop–Memory Uneven Mappings and Architectures for DNN Accelerators	879
<i>Xiaodong Liu, Zhihui Wang, Xiangcong Kong, Weixin Zhou, Xiaofei Xia, Yuqi Jiang</i>	
MARCO: Hardware-Aware Neural Architecture Search for Edge Devices with Multi-Agent Reinforcement Learning and Conformal Filtering	887
<i>Arya Fayyazi, Mehdi Kamal, Massoud Pedram</i>	
NetTLM-DSE: Design Space Exploration for DNN Layer-Pipeline Spatial Mappings.....	894
<i>Shinyoung Kim, Junsu Heo, Hyeseong Shin, Jaesuk Lee, Sungkyung Park, Chester Sungchung Park</i>	
BalanceGS: Algorithm-System Co-design for Efficient 3D Gaussian Splatting Training on GPU.....	903
<i>Junyi Wu, Jiaming Xu, Jinhao Li, Yongkang Zhou, Jiayi Pan, Xingyang Li, Guohao Dai</i>	
LumiLock: LUT-based Multi-Key Logic Locking	910
<i>Tsunato Nakai, Takuya Higashi</i>	
SCPrompt: Semantic Compression and Prompt-Guided LLM Reasoning for RTL Trojan Detection	917
<i>Jiaji He, Jiansheng Chen, Fei Zhao, Yaohua Wang, Yongqiang Lyu</i>	
Can Large Language Models Unlock Logic Locking?.....	924
<i>Takuya Higashi, Tsunato Nakai</i>	
GALA: An Explainable GNN-based Approach for Enhancing Oracle-Less Logic Locking Attacks Using Functional and Behavioral Features.....	931
<i>Yeganeh Aghamohammadi, Henry Jin, Amin Rezaei</i>	
SAND: A Self-supervised and Adaptive NAS-Driven Framework for Hardware Trojan Detection	938
<i>Zhixin Pan, Ziyu Shu, Linh Nguyen, Amberbir Alemayoh</i>	
An Effective Placement Framework for Designs with Half-Row-Extended Cells	945
<i>Po-Yi Wu, Tzu-Sheng Hung, Wai-Kei Mak, Ting-Chi Wang</i>	
DUALPlace: Reinforcement Learning based Mixed-size Placement with Multi Modal Cross Attention.....	952
<i>Yue Wu, Hu Liu, Jiayi Ding, Junwei Li, Tieming Han, Xiaoyan Yang</i>	
SMT-Based Optimal Transistor Folding and Placement for Standard Cell Layout Generation	959
<i>Junghyun Yoon, Heechun Park</i>	
Synthesis of CFET Standard Cells Utilizing Backside Interconnects Towards Improving Pin Accessibility	966
<i>Hyunbum Park, Taewhan Kim</i>	
Standard Cell Layout Synthesis for Dual-Sided 3D-Stacked Transistors.....	973
<i>Kairong Guo, Haoran Lu, Rui Guo, Jiarui Wang, Chunyuan Zhao, Heng Wu, Runsheng Wang, Yibo Lin</i>	
pHNSW: PCA-Based Filtering to Accelerate HNSW Approximate Nearest Neighbor Search.....	980
<i>Zheng Li, Guangyi Zeng, Paul Delestrac, Enyi Yao, Simei Yang</i>	

Boosting Scalability and Performance: Macro Placement for Flexible 3D-Stacked ML Accelerators.....	986
<i>Jiawei Hu, Canlin Zhang, Pruek Vanna-Iampikul, Tushar Krishna, Sung Kyu Lim</i>	
AutoCT: Hybrid Compressor Tree Optimization via Reinforcement Learning with Graph Modeling	993
<i>Shangshang Yao, Kunlong Li, Li Shen, Ruoxi Wang, Jingyu Liu</i>	
Chiplet-NAS: Chiplet-aware Neural Architecture Search for Efficient AI Inference on 2.5D Integration	999
<i>Cheng Guo, Pragnya Sudershan Nalla, Nikhil Kumar Cherukuri, Rui Xue, Sachin S. Sapatnekar, Chaitali Chakrabarti, Yu Cao, Jeff Zhang</i>	
MoBiLE: Efficient Mixture-of-Experts Inference on Consumer GPU with Mixture of Big Little Experts.....	1006
<i>Yushu Zhao, Yubin Qin, Yang Wang, Xiaolong Yang, Huiming Han, Shaojun Wei, Yang Hu, Shouyi Yin</i>	
StabiFreeze: Early Stopping for Training Binary Neural Networks via Internal Dynamics Stabilization.....	1013
<i>Jisu Kim, Tae-Hwan Kim</i>	
Constrained NAS via Symbolic Expressions in Declarative Hierarchical Search Spaces.....	1020
<i>Moritz Reiber, Christoph Gerum, Oliver Bringmann</i>	
HERO: Hardware-Efficient RL-based Optimization Framework for NeRF Quantization	1028
<i>Yipu Zhang, Chaofang Ma, Jinming Ge, Lin Jiang, Jiang Xu, Wei Zhang</i>	
Gundam: A Generalized Unified Design and Analysis Model for Matrix Multiplication on Edge	1035
<i>Quan Cheng, Haoyuan Li, Weirong Dong, Mingqiang Huang, Longyang Lin, Masanori Hashimoto</i>	
TTI: An Instruction Set Supporting Priority-Inversion-Free Time-Triggered Preemptive Scheduling in Real-Time Embedded Systems.....	1042
<i>Yinkang Gao, Yixuan Zhu, Bo Zhang, Haoyuan Ren, Cheng Tang, Xi Li</i>	
DRAPP: An end-to-end Latency Evaluation Tool for Containerized ROS Applications	1049
<i>Sicheng Guan, Haiyang Wang, Jinghao Sun, Qingxu Deng</i>	
LaPOD: Latency Prediction for Real-Time LiDAR Object Detection	1056
<i>Wenjing Xie, Tianchi Ren, Jen-Ming Wu, Chun Jason Xue, Nan Guan</i>	
Flowmap Overapproximation for Linear Time-Varying Stochastic Systems	1063
<i>Xin Chen</i>	
Timing-Aware Optimization of Die-Level Routing and TDM Assignment for Multi-FPGA Systems.....	1070
<i>Yijun Chen, Haoyuan Li, Chunyan Pei, Jianwang Zhai, Kang Zhao, Wenjian Yu</i>	
Φ -BO: Physics-Informed Bayesian Optimization for Multi-Port Decoupling Capacitor Placement in 2.5-D Chiplets	1077
<i>Quansen Wang, Yuchuan Lin, Zhuohua Liu, Wei Xing, Wei Zhang, Ning Xu, Yuanqing Cheng</i>	
Automated Parameter Tuning for Multi-FPGA Partitioning: A Preference-Guided Approach.....	1084
<i>Yutao Dai, Shengbo Tong, Chunyan Pei, Zhuohua Liu, Wei Xing, Yi Liu, Rui Wang, Wenjian Yu</i>	
Advancing General Sparse Linear-Equation Solvers via Nested-Dissection-Based Parallel Scheduling and Randomized Linear Algebra	1091
<i>Wenjian Yu, Jiawen Cheng, Baiyu Chen</i>	

IR Drop-Aware ECO: A Fast Approach to Minimize Layout and Timing Disturbance.....	1097
<i>Jingchao Hu, Yibo Lin, Hao Yu, Quan Chen, Zhou Jin, Cheng Zhuo</i>	
iPCL: Pre-training for Chip Layout (Invited Paper).....	1103
<i>Xingquan Li, Weiguo Li, Xinhua Lai, Junfeng Liu, Rui Wang</i>	
NEURAL: An Elastic Neuromorphic Architecture with Hybrid Data-Event Execution and On-the-fly Attention Dataflow	1110
<i>Yuehai Chen, Farhad Merchant</i>	
Mask-based Meta-Learning for Stuck-at Faults Tolerance in ReRAM Computing Systems	1117
<i>Zhan Shen, Yiqi Zhou, Tianyi Xu, Shan Shen, Zhen Mei, Daying Sun, Zichen Zhang</i>	
LMESN: A Leakage-Driven MOSFET Reservoir for Scalable and Ultra-Low-Power Temporal Inference	1124
<i>Masami Utsunomiya, Ryuko Seki, Quan Cheng, Weirong Dong, Feng Liang, Takashi Sato</i>	
BIHDC: A Retractable Fully-Binary Hyperdimensional Computing Accelerator for Edge FPGAs	1131
<i>Changzhen Han, Ke Chen, Bi Wu, Chenggang Yan, Weiqiang Liu</i>	
MOTIF-RF: Multi-template On-chip Transformer Synthesis Incorporating Frequency-domain Self-transfer Learning for RFIC Design Automation	1138
<i>Houbo He, Yizhou Xu, Lei Xia, Yaolong Hu, Fan Cai, Taiyun Chi</i>	
Efficient RF Passive Components Modeling with Bayesian Online Learning and Uncertainty Aware Sampling.....	1145
<i>Huifan Zhang, Pingqiang Zhou</i>	
PrometheusFree: Concurrent Detection of Laser Fault Injection Attacks in Optical Neural Networks	1152
<i>Kota Nishida, Yoshihiro Midoh, Noriyuki Miura, Satoshi Kawakami, Alex Orailoglu, Jun Shiomi</i>	
BEAM: Bidirectional MEEF-Driven Mask Optimization for Curvilinear Photonic Design.....	1160
<i>Xiaoxiao Liang, Yang Luo, Bei Yu, Yuzhe Ma</i>	
OAH-CIM: Outlier-Aware Hybrid RRAM-SRAM CIM Accelerator with Variation-Robust Sparsity	1167
<i>Zhiwei Zhou, Tong Hu, Han Bao, Houji Zhou, Yuyang Fu, Jiancong Li, Jia Chen, Yi Li, Xiangshui Miao</i>	
FlexMem: High-Parallel Near-Memory Architecture for Flexible Dataflow in Fully Homomorphic Encryption	1174
<i>Shangyi Shi, Husheng Han, Jianan Mu, Xinyao Zheng, Ling Liang, Hang Lu, Zidong Du, Xiaowei Li, Xing Hu</i>	
Data Flow-Aware Weight Remapping for Efficient Fault Tolerance in ReRAM-Based Accelerators.....	1181
<i>Hyeonsu Bang, Kang Eun Jeon, Jong Hwan Ko</i>	
MIREDO: MIP-Driven Resource-Efficient Dataflow Optimization for Computing-in-Memory Accelerator	1188
<i>Xiaolin He, Cenlin Duan, Yingjie Qi, Xiao May, Jianlei Yang</i>	
Hy2S-CIM: Hybrid-Cache-LUT FP/INT-CIM with 2-Stage Alignment and Area-efficient LUT for High Precision Vision AI Tasks	1195
<i>Xiang Li, Wenbin Jia, Sheng Zhang, Yongpan Liu</i>	
DOME: A Domain-Orchestrated Multi-GPU Optical Network for Rack-Scale Systems.....	1202
<i>Chongyi Yang, Peiyu Chen, Bohan Hu, Yinyi Liu, Wei Zhang, Jiang Xu</i>	

H2 NoC: A Hybrid NoC Architecture for FPGAs with Hardened Interconnects	1209
<i>Jinhyeong Park, Younggil Jeong, Jeongwoo Park</i>	
CommILP: Synthesizing Communication Infrastructure for Domain Computing Platforms.....	1216
<i>Qucheng Jiang, Jacob Ginesin, Oscar Kellner, Tianrui Ma, Gunar Schirner</i>	
ReadyPower: A Reliable, Interpretable, and Handy Architectural Power Model Based on Analytical Framework.....	1223
<i>Qijun Zhang, Shang Liu, Yao Lu, Mengming Li, Zhiyao Xie</i>	
A RISC-V CHERI VP: Enabling System-Level Evaluation of the Capability-Based CHERI Architecture	1230
<i>Manfred Schlagl, Andreas Hinterdorfer, Daniel Grose</i>	
EP-HDC: Hyperdimensional Computing with Encrypted Parameters for High-Throughput Privacy-Preserving Inference.....	1237
<i>Jaewoo Park, Chenghao Quan, Jongeun Lee</i>	
RTL Verification for Secure Speculation Using Cascaded Two-Phase Information Flow Tracking.....	1244
<i>Yuhao Liu, Ying Li, Yinhao Zhou, Chang Guo, Zhenfeng Li, Yanyun Lu</i>	
Exploiting Feature-driven Approximation to Preserve Privacy in Machine Learning based Health Monitoring Systems	1252
<i>Nishanth Chennagouni, Wei Lu, Qiaoyan Yu</i>	
Safeguarding Neural Network IPs from Scan Chain based Model Extraction Attacks	1259
<i>E Bhawani Eswar Reddy, Gundameedi Sai Ram Mohan, Sukanta Bhattacharjee, Chandan Karfa</i>	
FedBit: Accelerating Privacy-Preserving Federated Learning via Bit-Interleaved Packing and Cross-Layer Co-Design	1266
<i>Xiangchen Meng, Yangdi Lyu</i>	
DPO-3D: Differentiable Power Delivery Network Optimization via Flexible Modeling for Routability and IR-Drop Tradeoff in Face-to-Face 3D ICs.....	1273
<i>Zhen Zhuang, Zheng Yang, Yuxuan Zhao, Jiawei Hu, Bei Yu, Sung Kyu Lim, Tsung-Yi Ho</i>	
Enhancing Pin Accessibility Through Pin Pattern Migration and Optimization Across Cell Boundaries.....	1280
<i>Hyunbae Seo, Sehyeon Chung, Taewhan Kim</i>	
Topological Optimization-Based Layer Assignment Method for Fan-Out Wafer-Level Packaging	1287
<i>Haoying Wu, Guanxian Zhu, Xu He, Mingyu Liu</i>	
Au-MEDAL: Adaptable Grid Router with Metal Edge Detection And Layer Integration	1294
<i>Andrew B. Kahng, Seokhyeong Kang, Sehyeon Kim, Jakang Lee, Dooseok Yoon</i>	
GPU-Accelerated Global Routing with Balanced Timing and Congestion Optimization	1301
<i>Jinghui Zhou, Fuxing Huang, Lixin Chen, Xinglin Zheng, Ziran Zhu</i>	
Fine-Grained Parallelization of FHE Workloads in Multi-GPU Systems.....	1308
<i>Homer Gamil, Michail Maniatakos</i>	
Exploiting the Irregular Input Sparsity in Systolic Array-based DNN Accelerators via Local Soft Pooling	1315
<i>Desheng Fu, Jingbo Jiang, Xiaomeng Wang, Jingyang Zhu, Xizi Chen, Chi-Ying Tsui</i>	

Input Reuse, Weight-Stationary Dataflow and Mapping Strategy for Depthwise Convolution in Computing-in-Memory Neural Network Accelerators	1322
<i>Chia-Chun Wang, Yu-Chih Tsai, Ren-Shuo Liu</i>	
A Unified Compute-In-Memory Framework for Multisensory Emotion Recognition	1329
<i>He Xiao, Yue Zhou, Dirui Xie, Qi Cheng, Xiaofang Hu, Zhengwu Liu, Ngai Wong</i>	
Scope: A Scalable Merged Pipeline Framework for Multi-Chip-Module NN Accelerators	1336
<i>Zongle Huang, Hongyang Jia, Kaiwei Zou, Yongpan Liu</i>	
UltraMalloc: Efficient FPGA-based Memory Allocation Framework Optimized for HBM	1343
<i>Yuwei Qu, Yiqing Mao, Yanxing Jin, Wai-Shing Luk, Lingli Wang</i>	
Viper: An ILP-Based Vectorization Framework for Fully Homomorphic Encryption	1350
<i>Weidong Yang, Xinmo Li, Xiangmin Guo, Jianfei Jiang, Naifeng Jing, Qin Wang, Zhigang Mao, Weiguang Sheng</i>	
WPU: A Pipelined WebAssembly Processing Unit for Embedded IoT Systems	1357
<i>Jinyeol Kim, Chaebin Lee, Jongwon Oh, Joungmin Park, Seung Eun Lee</i>	
SARA: A Stall-Aware Memory Allocation Strategy for Mixed-Criticality Systems	1364
<i>Meng-Chia Lee, Wen Sheng Lim, Yuan-Hao Chang, Tei-Wei Kuo</i>	
Zone-aware metadata placement in B-tree filesystem	1372
<i>Ming-Feng Wei, Yun-Chih Chen, Yuan-Hao Chang, Tei-Wei Kuo</i>	
Full-Chip Thermal Map Estimation by Multimodal Data Fusion via Denoising Diffusion	1379
<i>Jing Li, Wei Xing, Yuquan Sun, Yuanqing Cheng</i>	
Graph Attention-Based Current Crowding Analysis at TSV Interfaces in 3D Power Delivery Networks	1386
<i>Zheng Yang, Zhen Zhuang, Tsung-Yi Ho, Sung Kyu Lim</i>	
Efficient Simulation of IC Packages with TEC Based on Adaptive Segmented Method and Spatially-Aware Thermal Neural Network	1393
<i>Longfei Li, Xinyue Li, Shunxiang Lan, Liang Chen, Haibao Chen, Min Tang</i>	
Domain Transformation and Decomposition Method for Composable Thermal Modeling and Simulation of Chiplet-Based 2.5D Integrated System	1400
<i>Shunxiang Lan, Min Tang</i>	
ThPA: Thermal Simulation for Advanced ICs	1407
<i>Bo-Wen Chen, Yong-Han Lin, Chien-Yu Lin, Yu-Min Lee</i>	
DRLPlace: A Deep Reinforcement Learning-based Irregular and High-Density Printed Circuit Board Placement Method	1414
<i>Lei Cai, Ke Cheng, Jixin Zhang, Haiyun Li, Zhiwei Ye</i>	
CNN-Assisted Low-Power Clock Tree Synthesis for 3D ICs	1421
<i>Chenbo Xi, Jindong Zhou, Pingqiang Zhou</i>	
DALI-PD: Diffusion-based Synthetic Layout Heatmap Generation for ML in Physical Design	1428
<i>Bing-Yue Wu, Vidya A. Chhabria</i>	
Partitioning-free 3D-IC Floorplanning	1435
<i>Shuo Ren, Zhen Zhuang, Rongliang Fu, Leilei Jin, Libo Shen, Bei Yu, Tsung-Yi Ho</i>	

A Heterogeneous Graph-based Gate Sizer Integrating Graph Attention Network and Transformer	1442
<i>Jinmo Ahn, Jinoh Cho, Jaemin Seo, Jaeseung Lee, Jakang Lee, Seokhyeong Kang</i>	
Platinum: Path-Adaptable LUT-Based Accelerator Tailored for Low-Bit Weight Matrix Multiplication	1449
<i>Haoxuan Shan, Cong Guo, Chiyue Wei, Feng Cheng, Junyao Zhang, Hai Helen Li, Yiran Chen</i>	
ROMA: A Read-Only-Memory-based Accelerator for QLoRA-based On-Device LLM.....	1456
<i>Wenqiang Wang, Yijia Zhang, Zikai Zhang, Guanting Huo, Hao Liang, Shijie Cao, Ningyi Xu</i>	
Noise-Agnostic One-Shot Training and Retraining for Robust DNN Inferencing on Analog Compute-in-Memory Systems.....	1463
<i>Ashish Reddy Bommana, Ben Feinberg, Patrick Xiao, Christopher Bennet, Matthew Marinella, Krishnendu Chakrabarty</i>	
Activation-Free Implicit Neural Representation via Finite-State-Machine Based Stochastic Computing.....	1470
<i>Xincheng Feng, Wenyong Zhou, Taiqiang Wu, Meng Li, Zhengwu Liu, Ngai Wong</i>	
dLLM-OPU: An FPGA Overlay Processor for Accelerated Diffusion Large Language Models	1476
<i>Yangbo Wei, Shaoqiang Lu, Junhong Qian, Lei He, Dongge Qin, Xiao Shi, Chen Wu, Linfeng Zhang</i>	

Author Index