

2026 Design, Automation & Test in Europe Conference (DATE 2026)

Verona, Italy

20-22 April 2026

Pages 1–676

IEEE Catalog Number:
ISBN:

CFP26162-POD
979-8-3315-4565-9



Copyright © 2026, EDAA

All Rights Reserved

Copyright and Reprint Permissions:

Abstracting is permitted with credit to the source. For permission to reprint or republish any portion of this publication, please contact the copyright holder.

*** This is a print representation of what appears in an electronic format. Some format issues inherent in the e-media version may also appear in this print version.

IEEE Catalog Number:	CFP26162-POD
ISBN (Print-On-Demand):	979-8-3315-4565-9
ISBN (Online):	978-3-9826-7411-7
ISSN:	1530-1591

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

TABLE OF CONTENTS

AERO: Adaptive and Efficient Runtime-Aware OTA Updates for Energy-Harvesting IoT	1
<i>Wei Wei, Jingye Xu, Sahidul Islam, Dakai Zhu, Chen Pan, Mimi Xie</i>	
Bespoke Co-processor for Energy-Efficient Health Monitoring on RISC-V-based Flexible Wearables	8
<i>Theofanis Vergos, Polykarpos Vergos, Mehdi B. Tahoori, Georgios Zervakis</i>	
ALIFE-BCI: An Adaptive Low-power Integrated Feature Extractor for Brain-Computer Interfaces	15
<i>Joe Saad, Ivan Miro-Panades, Adrian Evans, Lorena Anghel</i>	
Efficient LLM Decoding on Ryzen AI NPUs	22
<i>Zhenyu Xu, Miaoxiang Yu, Jillian Cai, Qing Yang, Tao Wei</i>	
EGO: Efficient Compression of Unstructured Sparse DNNs for Compute-in-Memory based on Graph Minimum-Cost Matching Optimization	29
<i>Teng Wan, Yu Cao, Huazhong Yang, Xueqing Li</i>	
Grin: HyperGNN Training Framework for Efficient Edge Inference via Hypergraph Restructuring	35
<i>Chaofang Ma, Lin Jiang, Zeyu Li, Xingyu Liu, Jiang Xu, Wei Zhang</i>	
A Scheduling Framework for Efficient MoE Inference on Edge GPU-NDP Systems	42
<i>Qi Wu, Chao Fang, Jiayuan Chen, Ye Lin, Yueqi Zhang, Yichuan Bai, Yuan Du, Li Du</i>	
Colored Huge Pages: A Hardware-Software Approach for Enhanced Isolation and Performance	49
<i>Georgios-Alexandros Kostas, Dimitris Gizopoulos, Vasileios Karakostas</i>	
An Open Source Design Exploration Tool for Battery and Coolant Configuration	56
<i>Francesco Tosoni, Yukai Chen, Massimo Poncino, Franco Fummi, Sara Vinco</i>	
Exploring Heterogeneity-Aware Optimizations for Resource Efficient Edge Recommendation	59
<i>Yerin Lee, Gyudong Kim, Eunjin Lee, Jeff Zhang, Young-Ho Gong, Young Geun Kim, Carole-Jean Wu</i>	
Glitch Propagation through Flip-Flops Endangers Masking Schemes: Why Time Separation Is Required	62
<i>Hasin Ishraq Reefat, Mohammad Ebrahimabadi, Sofiane Takarabt, Sylvain Guilley, Naghmeh Karimi</i>	
Distilling Graph Reasoning into Lightweight CNNs for Near-Sensor Point Cloud Corruption Detection	69
<i>Grafika Jati, Martin Molan, Francesco Barchi, Andrea Bartolini, Giuseppe Mercurio, Andrea Acquaviva</i>	
FHEx: Transforming Generic Compute Chips into Secure FHE Engines via a Hardware-software Co-designed Framework	76
<i>Yibo Du, Ying Wang, Mengdi Wang, Cangyuan Li, Lian Liu, Hui Li, Kai Zhang, Yinhe Han</i>	
Towards Trustworthy LLM-Based Assertion Generation: A Data Augmentation Framework with Formal Check Approach	83
<i>Qingchen Zhai, Hao Yu, Chen Bai, Charles Young, Frank Qu, Dezhi Ran, Yuan Xie, Tao Xie</i>	
Automated Self-Explanation of Expected versus Perceived Behavior for Interacting Digital Systems	90
<i>Mohammad Alkhiyami, Gianluca Martino, Goerschwin Fey</i>	
Exploring a Resource-Efficient NTT FPGA Accelerator for Fully Homomorphic Encryption	97
<i>Valentino Guerrini, Giuseppe Sorrentino, Davide Conficconi</i>	

ML-DSA-OSH: An Efficient, Open-Source Hardware Implementation of ML-DSA	100
<i>Quinten Norga, Suparna Kundu, Ingrid Verbauwhede</i>	
Postponing the Glitches is not Enough A Critical Analysis of the DATE 2024 E-ISW Masking Scheme	103
<i>Amir Moradi</i>	
From Trigger to Impact: Knowledge-Graph Reasoning and Risk-Aware Classification for Hardware Trojan Detection	106
<i>Yang Zhang, Xing Hu, Xiaowen Chen, Huan Guo, Zhenyu Zhao, Sheng Liu</i>	
Eunomia: Preemption-based and QoS-aware Core Allocation in Oversubscribed Cloud	108
<i>Yunda Guo, Puqing Wu, Haoqiong Bian, Yunpeng Chai, Yao Shen, Haoyu Yang, Qing Liu, Zhengbin Huang Le Yue, Yi Yang</i>	
Breaking Standard Cell Margin Constraints for Area-Efficient VLSI Design	115
<i>Junghyun Yoon, Jooyeon Jeong, Heechun Park</i>	
Accelerating Detailed Routing Convergence through Offline Reinforcement Learning	122
<i>Afsara Khan, Austin Rovinski</i>	
Communication-Aware Hybrid Parallelism Mapping for Low-Cost MCM-based DNN Accelerators	129
<i>Jicheon Kim, Chunmyung Park, Xuan Truong Nguyen, Hyuk-Jae Lee</i>	
A Reusable Methodology for High-Performance Interconnects using a Standard-Cell based Asynchronous NoC Router	136
<i>Chonghui Zhang, Yizhe Hu, Yi Kang</i>	
MONET: A Mixture-of-Experts Accelerator with a Multicast-Optimized Two-Tier Network-on-Chip	143
<i>Siqin Liu, Maya Roediger, Avinash Karanth</i>	
Efficient Throughput Analysis of Synchronous Dataflow Graphs via Parametric Shortest Path	150
<i>Zhengzheng Tian, Mingze Ma, Jian Hou</i>	
Unified Class and Domain Incremental Learning with Mixture of Experts for Indoor Localization	157
<i>Akhil Singampalli, Sudeep Pasricha</i>	
Compression Space Search: RL-Based Combinational Compression for Neural Networks	164
<i>Yingtao Shen, Yinchun Ni, Jiace Zhu, Jie Zhao, An Zou</i>	
TrainDeploy: Hardware-Accelerated Parameter-Efficient Fine-Tuning of Small Transformer Models at the Extreme Edge	171
<i>Run Wang, Victor J.B. Jung, Philip Wiese, Francesco Conti, Alessio Burrello, Luca Benini</i>	
QSLM: A Performance- and Memory-aware Quantization Framework with Tiered Search Strategy for Spike-driven Language Models	178
<i>Rachmad Vidya Wicaksana Putra, Pasindu Wickramasinghe, Muhammad Shafique</i>	
FiCABU: A Fisher-Based, Context-Adaptive Machine Unlearning Processor for Edge AI	185
<i>Eun-Su Cho, Jongin Choi, Jeongmin Jin, Jae-Jin Lee, Woojoo Lee</i>	
HILAL: Hessian-Informed Layer Allocation for Heterogeneous Analog–Digital Inference	192
<i>Aniss Bessalah, Hatem Mohamed Abdelmoumen, Karima Benatchba, Hadjer Benmeziane</i>	

PatchBlock: A Lightweight Defense Against Adversarial Patches for Embedded EdgeAI Devices	199
<i>Nandish Chattopadhyay, Abdul Basit, Amira Guesmi, Muhammad Abdullah Hanif, Bassem Ouni Muhammad Shafique</i>	
Efficient CNN Inference on Ultra-Low-Power MCUs via Saturation-Aware Convolution	206
<i>Shiming Li, Luca Mottola, Yuan Yao, Stefanos Kaxiras</i>	
Lupin: Spatial Resource Stealing with Outlier-First Encoding for Mixed-Precision LLM Acceleration	209
<i>Taein Kim, Sukhyun Han, Seongwook Kim, Gwangeun Byeon, Jungmin Lee, Seokin Hong</i>	
MAPLE: Modality-Aware Projection-free LiDARCamera Fusion for 3D Vehicular Object Detection	212
<i>Abhishek Balasubramaniam, Sudeep Pasricha</i>	
3D Integration of Hybrid IGZO/Si and IGZO eDRAMs for High-Density/High-Performance On-Chip Memory	215
<i>Munhyeon Kim, Sukhyun Choi, Yulhwa Kim, Jae-Joon Kim</i>	
An Operator-Circuit Co-design Digital SOT-MRAM Computing-in-Memory Accelerator with Double Bit Density and Full-Utilized Bandwidth/Throughput	222
<i>Tianshuo Bai, Jingcheng Gu, Lehao Tan, Wenten Yi, Haolin Ge, Han Zhang, Zhenyu Xue, He Zhang, Na Lei Biao Pan</i>	
FSDB: A Folded-Store Dynamic-Broaden Hybrid Compute-in-ROM/SRAM Architecture for Deploying Large-Scale DNNs On-Chip	229
<i>Tianyi Yu, Teng Yi, Huazhong Yang, Xueqing Li</i>	
The Munich Quantum Software Company: Developing Production-ready Quantum Computing Software	236
<i>Robert Wille, Marcel Walter, Simon Hofmann, Patrick Hopf, Marc Messing, Lukas Burgholzer</i>	
Artemis: Co-Simulation of Power Microgrids and Energy-Aware Cloud Data Centers	237
<i>Mattia Tibaldi, Sara Vinco, Christian Pilato</i>	
COFFEE: A Carbon-Modeling and Optimization Framework for HZO-based FeFET eNVMS	244
<i>Hongbang Wu, Xuesi Chen, Shubham Jadhav, Amit Lal, Lillian Pentecost, Udit Gupta</i>	
Reshaping Bayesian Optimization for Design Space Optimization Towards Accurate and Irredundant Evaluation in EDA Tool Parameter Exploration	251
<i>Chanhee Jeon, Taewhan Kim</i>	
EDA Flow Matters: Stage-Aware Parameter Optimization of Tool Chain	258
<i>Xinheng Li, Donger Luo, Peng Xu, Ziyang Yu, Qi Sun, Tinghuan Chen, Bei Yu, Hao Geng</i>	
FAQNAS: FLOPs-aware Hybrid Quantum Neural Architecture Search using Genetic Algorithm	265
<i>Muhammad Kashif, Shaf Khalid, Alberto Marchisio, Nouhaila Innan, Muhammad Shafique</i>	
Unified Pauli-Rotation Synthesis for Relieving CX-Count Overhead in Tableau-Based Quantum Circuit Optimization Flow	272
<i>Yi-Hsiang Kuo, Hsiang-Chun Yang, Hsin-Yu Chen, Chung-Yang Ric Huang</i>	
ChipLight: Cross-Layer Optimization of Chiplet Design with Optical Interconnects for LLM Training	279
<i>Kangbo Bai, Zhantong Zhu, Yifan Ding, Tianyu Jia</i>	
Cyber-Physical System Design Space Exploration for Affordable Precision Agriculture	286
<i>Pawan Kumar, Hokeun Kim</i>	

Design and Optimization of Solar-Powered Embedded Systems with Uppaal Stratego	293
<i>Ismael Samaye, Abdoulaye Gamatié</i>	
EMaper: Cross-level Electromigration Aware Placement and Routing EDA Workflow for Interconnects	
Hotspot Prediction and Mitigation	296
<i>Chenglin Ye, Yuze Lu, Yizhan Liu, Ligong Zhang, Jinghan Xu, Fei Liu, Yibo Lin, Zheng Zhou, Xiaoyan Liu</i>	
ETLA-3D: Equivalent Thin Layer Aggregation based Thermal FEM for Hybrid Bonding F2F 3D ICs	303
<i>Chenghan Wang, Zhen Zhuang, Kai Zhu, Darong Huang, Luis Costero, Rongmei Chen, David Atienza</i>	
<i>Tsung-Yi Ho</i>	
Consolidating ML-driven Early IR-drop Mitigation for Fast and Reliable IR-drop Closure	310
<i>Munwon Lee, Chanhee Jeon, Taewhan Kim</i>	
FALCON-3D: Full-Chip Analytical Thermal Simulation with Lateral CONvection for 3D-Stacked ICs	317
<i>Tsung-Lin Lu, Yu-Min Lee, Pei-Yu Huang, Ching-Hsiang Wang</i>	
3D-ICE 4.0: Accurate and efficient thermal modeling for 2.5D/3D heterogeneous chiplet systems	324
<i>Kai Zhu, Darong Huang, Luis Costero, David Atienza</i>	
Sensor Placement and Transformer-Based Thermal Map Generation for Reusable Interposers	331
<i>Aristotelis Tsekouras, Theodoros Papavasileiou, Panagiotis C. Petrantonakis, Georgios Keramidas</i>	
<i>Vasilis F. Pavlidis</i>	
Enabling Cross-Design Power Trace Prediction with GNNs for Gate-Level Netlists	338
<i>Shih-Chun Lin, Yung-Chih Chen, Bo-Hao Huang</i>	
FastRW: An Efficient Random Walk Method for Steady-State Thermal Analysis	345
<i>Zixiao Wang, Tianshu Hou, Chenghan Wang, Zhen Zhuang, Tsung-Yi Ho, Farzan Farnia, Bei Yu</i>	
Reinforcement Learning for Hybrid Bonding Terminal Legalization in 3D ICs	348
<i>Wanqi Ren, Chengrui Gao, Yunqi Shi, Mingzhou Fan, Siyuan Xu, Ke Xue, Chenjian Ding, Mingxuan Yuan</i>	
<i>Chao Qian</i>	
Think with Self-Decoupling and Self-Verification: Automated RTL Design with Backtrack-ToT	351
<i>Zhiteng Chao, Yonghao Wang, Xinyu Zhang, Jiabin Zhou, Tenghui Hua, Husheng Han, Tianmeng Yang</i>	
<i>Jianan Mu, Bei Yu, Rui Zhang, Jing Ye, Huawei Li</i>	
HiM: An Autonomous Hardware Accelerator for Solving Boolean Satisfiability Problem with a	
Heuristic-in-Macro Engine	358
<i>Shin Han, Minhyeok Jeong, Yoonmyung Lee</i>	
MISP-Net: Significantly Reducing Transient Backward Steppings via Novel Multi-step Irregular	
Sequence Prediction	365
<i>Yichao Dong, Dan Niu, Chao Wang, Zhenya Zhou, Zhou Jin, Changyin Sun</i>	
Accurate Analytical Modeling for NoCs with Hybrid Arbitration under High Traffic Injection	372
<i>Rahul Tripathy, Mohammad Majharul Islam, Riad Akram, Raid Ayoub, Sumit K. Mandal</i>	
A Parallel Mixed-Precision GMRES-IR Solver for Ill-Conditioned Equations in Device Simulation	379
<i>Jiawen Cheng, Yuanyuan Yang, Ding Gong, Wenjian Yu</i>	
PALM: Program Analysis and LLM Methods for Crafting SystemVerilog Assertions	386
<i>Raheel Afsharmazayehani, Benjamin Tan</i>	

Contract-Based Architecture Exploration of Cyber-Physical Systems via Satisfiability Modulo Convex Programming	393
<i>Yifeng Xiao, Pierluigi Nuzzo</i>	
DyNAMoE: Dynamic Reconfigurable NoC-based Accelerator for Mixture-of-Expert Models	400
<i>Mohit Upadhyay, Li-Shiuan Peh</i>	
Dynamic Rank-Aware Aggregation with Graph Contrastive Learning for Federated Foundation Model Fine-Tuning	403
<i>Zhao Yang, Xunyun Qiu, Hua Cui</i>	
FedTPA: Tackling Data Heterogeneity with Adaptive Parameter Allocation in Federated Instruction Tuning	410
<i>Yixuan Chen, Jinghui Zhang, Ding Ding</i>	
Shakan: Training-Inference co-design for Oblique Random Forests on Embedded Devices	417
<i>Alessandro Annechini, Alessandro Verosimile, Marco D. Santambrogio</i>	
Dynamic Neural Thresholding on Mixed-Signal Neuromorphic Processors Enabled by Integrated Learning and Hardware Design	424
<i>Kyuseung Han, Kwang-Il Oh, Sukho Lee, Hyeonuk Jang, Jae-Jin Lee, Sooyoung Jang</i>	
Vmxdotp: A RISC-V Vector ISA Extension for Efficient Microscaling (MX) Format Acceleration	431
<i>Max Wipfli, Gamze İslamoğlu, Navaneeth Kunhi Purayil, Angelo Garofalo, Luca Benini</i>	
Improving Reliability in Quantized Graph Neural Networks with Node-Wise Entropy-driven Temperature Scaling	438
<i>Hadi Mousanejad Jeddi, Jose Nunez-Yanez</i>	
A Self-Supervised Neuromorphic Processor Using High-Dimensional Representations for Cognitive Map Navigation	445
<i>Anqin Xiao, Luyu Yang, Yuhan He, Hengtan Zhang, Ziyi Yang, Lirong Zheng, Zhuo Zou</i>	
Parallel-SA: Point Cloud Processing Acceleration via Parallel Set Abstraction	452
<i>Dongdong Tang, Weilan Wang, Yu Mao, Wenjing Xie, Nan Guan, Tei-Wei Kuo, Chun Jason Xue</i>	
SONIC: Smart Optimization for Neural-Integrated CMP with Timing-Aware Fills	459
<i>Jiajun Tan, Qichao Ma, Yiming Du, Yiming Gan, Ling Lang, Yibo Lin, Ming Zhu, Zongwei Wang Yimao Cai</i>	
MPM-LLM4DSE: Reaching the Pareto Frontier in HLS with Multimodal Learning and LLM-Driven Exploration	462
<i>Lei Xu, Shanshan Wang, Chenglong Xiao</i>	
DAPO: Design Structure-Aware Pass Ordering for HLS via Contrastive and Reinforcement Learning	469
<i>Jinming Ge, Linfeng Du, Likith Anaparty, Shangkun Li, Tingyuan Liang, Afzal Ahmad, Vivek Chaturvedi Sharad Sinha, Zhiyao Xie, Jiang Xu, Wei Zhang</i>	
AutoShrink: Adaptive Search Space Shrinkage for Large-Scale Pareto Optimization of HLS Designs	476
<i>Yingxin Zeng, Binghao Cheng, Jianwang Zhai, Kang Zhao, Zhe Lin</i>	
Area Efficient Speculative Loop Pipelining for High-Level Synthesis	483
<i>Dylan Leothaud, Simon Rokicki, Steven Derrien, Isabelle Puaut</i>	

Anchor-and-Adapt: HLS QoR Prediction using Ground-Truth Seeding and Few-Shot Fine-Tuning	490
<i>Gabriel C. Tavares, Heitor C. De Andrade, Fábio P. Itturriet, Gabriel L. Nazar</i>	
Bench4HLS: End-to-End Evaluation of LLMs in High-Level Synthesis Code Generation	497
<i>M Zafir Sadik Khan, Kimia Azar, Hadi Kamali</i>	
OpenACM: An Open-Source SRAM-Based Approximate CiM Compiler	504
<i>Yiqi Zhou, JunHao Ma, Xingyang Li, Yule Sheng, Yue Yuan, Yikai Wang, Bochang Wang, Yiheng Wu Shan Shen, Wei Xing, Daying Sun, Li Li, Zhiqiang Xiao</i>	
RAPID-Graph: Recursive All-Pairs Shortest Paths Using Processing-in-Memory for Dynamic Programming on Graphs	511
<i>Yanru Chen, Zheyu Li, Keming Fan, Runyang Tian, John Hsu, Weihong Xu, Minxuan Zhou, Tajana Rosing</i>	
PiCoSNN: Partially Incoherent Configurable Optical Computing Architecture for SNN Acceleration	518
<i>Bowen Duan, Zhenhua Zhu, Zhengyang Duan, Huazhong Yang, Yuan Xie, Yu Wang</i>	
LAMP: An Adaptive Near-Memory Processing System for High-Performance Long-Read Mapping	525
<i>Jo-Ling Huang, Liang-Chi Chen, Chien-Chung Ho, Yuan-Hao Chang</i>	
MD-SNN: Membrane Potential-aware Distillation on Quantized Spiking Neural Network	532
<i>Donghyun Lee, Abhishek Moitra, Youngeun Kim, Ruokai Yin, Priyadarshini Panda</i>	
A Digital Neural Array IC for Real-Time Neural Network Replication from Spike Activities	539
<i>Donghyun Park, Hajung Mun, Minhyeok Jeong, Dahee Kang, Jongmin Lee, Yoonmyung Lee</i>	
GMaC: NvCIM Architecture for Parallel Point-based Point Cloud Acceleration via Geometric Mapping and Address-Index Computation	546
<i>Yi Gao, Zongwei Wang, Ling Liang, Yimao Cai</i>	
FSR-GeMM: A Scalable FSR-Parallel Photonic Accelerator for Real-Valued GeMM Computing	553
<i>Peiyu Chen, Yinyi Liu, Minhang Xu, Chongyi Yang, Xiaohan Jiang, Wei Zhang, Jiang Xu</i>	
RETRO: Mitigating Power Side-Channel Attacks with Reconfigurable RFET-based Ring Oscillators	560
<i>Nima Kavand, Tushar Niranjana, Armin Darjani, Akash Kumar</i>	
C-STEP: Compute-Efficient Spiking Transformers with Temporal Exit and Early-Guided Pruning	567
<i>Kyunchul Lee, Jongsun Park</i>	
Agentic AI for Digital Wellness: Challenges and Architectural Perspectives for Smart Home Care	570
<i>Luigi Capogrosso, Francesco Biondani, Francesca Bigardi, Stefano Cordibella, Giovanni Perbellini Walter Vendraminetti, Franco Fummi</i>	
LaMoS: Enabling Efficient Large Number Modular Multiplication through SRAM-based CiM Acceleration	573
<i>Haomin Li, Fangxin Liu, Chenyang Guan, Zongwu Wang, Li Jiang, Haibing Guan</i>	
MARS: A General GPU Optimization Framework for Merkle-Tree-Enabled Cryptography	580
<i>Yaoyun Zhou, Qian Wang</i>	
Attacking and Securing Hybrid Homomorphic Encryption Against Power Analysis	587
<i>Aikata Aikata, Maciej Czaprynski, Nedžma Mušović, Emira Salkić, Sujoy Sinha Roy</i>	
UniSEC: A Unified Security Evaluation Framework for Secure Cache Architectures	594
<i>Pratik Shrestha, Achim D. Brucker, M. Khurram Bhatti</i>	

FHEIns: Fully Homomorphic Encryption Acceleration for Large Data Applications with In-Storage Processing	601
<i>Xuan Wang, Tianqi Zhang, Keming Fan, Augusto Vega, Minxuan Zhou, Tajana Rosing</i>	
Make it Darker: A Gray Code Popcounter to Protect BNN CIM Against Power Attacks	608
<i>Fouwad Jamil Mir, Asmae El Arrassi, Abdullah Aljuffri, Said Hamdioui, Mottaqiallah Taouil</i>	
Efficient Federated Learning with Low-Rank Updates under Homomorphic Encryption	615
<i>Mohamed Aboelenien Ahmed, Mohamed Alsharkawy, Hassan Nassar, Heba Khdr Jeferson Gonzalez-Gomez, Jörg Henkel</i>	
Scytale: A Compiler Framework for Accelerating TFHE with Circuit Bootstrapping	622
<i>Rostin Shokri, Nektarios Georgios Tsoutsos</i>	
SPOILER-GUARD: Gating Latency Effects of Memory Accesses through Randomized Dependency Prediction	629
<i>Gayathri Subramanian, Girinath P, Nitya Ranganathan, Kamakoti Veezhinathan Gopalakrishnan Srinivasan</i>	
CIRCE CROSS Integrated RISC-V Cryptographic Extension	632
<i>Alessandra Dolmeta, Valeria Piscopo, Maurizio Martina, Guido Masera</i>	
Quantum Circuit Compilation for Superconducting Bus-Resonator Architectures	635
<i>Patrick Hopf, Lukas Burgholzer, Robert Wille</i>	
Alternating ZX Circuit Extraction for Hardware-Adaptive Compilation	642
<i>Ludwig Schmid, Korbinian Staudacher, Robert Wille</i>	
SurgeQ: A Hybrid Framework for Ultra-Fast Quantum Processor Design and Crosstalk-Aware Circuit Execution	649
<i>Xinxuan Chen, Hongxiang Zhu, Zhaohui Yang, Zhaofeng Su, Jianxin Chen, Feng Wu, Huihai Zhao</i>	
Code Division Multiplexing based Readout Scheme for Spin Qubits	656
<i>Jean-Baptiste Casanova, Quentin Schmidt, Baptiste Jadot, Brian Martinez, Xavier Jehl, Franck Badets Yvain Thonnart</i>	
Quantum Circuit Synthesis Based on LimTDD	663
<i>Xin Hong, Chenjian Li, Aochu Dai, Runhong He, Shenggang Ying</i>	
Quantum Hardware-Efficient Selection of Auxiliary Variables for QUBO Formulations	670
<i>Damian Rovara, Lukas Burgholzer, Robert Wille</i>	
Optimal Compilation of Syndrome Extraction Circuits for General Quantum LDPC Codes	677
<i>Kai Zhang, Dingchao Gao, Zhaohui Yang, Runshi Zhou, Fangming Liu, Zhengfeng Ji, Jianxin Chen</i>	
Ultra-Low Logical Depth Fault-Tolerant Quantum Circuit Synthesis via Lattice Surgery	684
<i>Chien-Tung Cherie Kuo, Cheng-En Tsai, Chung-Yang Ric Huang</i>	
Efficient Image Reconstruction Architecture for Neutral Atom Quantum Computing	691
<i>Jonas Winklmann, Yian Yu, Xiaorang Guo, Korbinian Staudacher, Martin Schulz</i>	
A Mathematical Exploration to Equivalence Checking of Quantum Circuits	694
<i>You-Cheng Lin, Yi-Ting Li, Wuqian Tang, Yung-Chih Chen, Chia-Chieh Chu, Chun-Yao Wang</i>	

From Forest to Tree: Prioritizing the Maximum Additional Delay in AQFP Circuit Design	697
<i>Yinuo Bai, Mingjia Fan, Tsung-Yi Ho, Zhou Jin</i>	
SCALER: A Stream-Aware Accelerator with Hierarchical Memory for Sparse LU Factorization on HBM FPGAs	700
<i>Xin Xu, Zhiying Zhu, Zishu Li, Dan Niu, Cheng Zhuo, Zhou Jin</i>	
Torrent : A Distributed DMA for Efficient and Flexible Point-to-Multipoint Data Movement	707
<i>Yunhao Deng, Fanchen Kong, Xiaoling Yi, Ryan Antonio, Marian Verhelst</i>	
KirbyMM: Outer-Product Based Matrix Multiplication on ARMv9 Processor	714
<i>Lanshu Huang, Han Huang, Zhiguang Chen, Yutong Lu</i>	
Focus Session: Hardware/Software Co-Design to Accelerate Generative AI Workloads on Heterogeneous Architectures	721
<i>Pratyush Dhingra, Vibhanshu Sharma, Janardhan Rao Doppa, Partha Pratim Pande</i>	
Focus Session: Accelerating Diffusion Models for Generative AI Applications with Silicon Photonics	728
<i>Tharini Suresh, Salma Afifi, Sudeep Pasricha</i>	
Focus Session: Hardware and Software Techniques for Accelerating Multimodal Foundation Models	735
<i>Muhammad Shafique, Abdul Basit, Muhammad Abdullah Hanif, Alberto Marchisio Rachmad Vidya Wicaksana Putra, Minghao Shao</i>	
Gohan: A Golden-Copy-Aided Platform Enabling Online Hybrid-Interactive Reliability Analysis	742
<i>Quan Cheng, Haoyuan Li, Wang Liao, Feng Liang, Longyang Lin, Masanori Hashimoto</i>	
SNIFFER: RL-based Vendor-Agnostic Test Case Generation for Triggering Long-Latency Behaviors	749
<i>Mingyu Pi, Michael Yun, Jaeseung Seok, Sangmin Kim, Sunghye Lee, Jinhwa Lee, Yoon Hyeok Lee</i>	
Explainable GNN-Driven Test Point Insertion on Uncontrollable I/Os	756
<i>Sung-Hyuk Cho, Tae-Min Park, Jeongyeol Lee, Jae-Youn Hong, Andreas Gerstlauer, Joon-Sung Yang</i>	
Design for testability using mixed-polarity flip-flops and latches	763
<i>Lorenzo Lagostina, Jordi Cortadella, Mario R. Casu, Luciano Lavagno</i>	
Adaptive Testing of Compute-in-Memory GANs Using Backpropagation-Guided Test Compaction	769
<i>Anurup Saha, Ashiqur Rasul, Thomas A. Walton, Amirali Aghazadeh, Abhijit Chatterjee</i>	
A Reinforcement Learning Framework for Good Die in Bad Neighborhood Analysis	776
<i>Mohammad Ershad Shaik, Abhishek Mishra, Nagarajan Kandasamy, Nur A. Toubia</i>	
Concurrent Fault Detection for Binary Neural Network Accelerators via On-Chip Voltage Monitoring	783
<i>Vincent Meyers, Mahboobe Sadeghipour Roodsari, Mehdi Tahoori</i>	
Noise-Aware Adaptive Sampling for Robust Diffusion Models on Analog Compute-in-Memory	790
<i>Yuannuo Feng, Wenyong Zhou, Yuexi Lv, Hanjie Liu, Guangyao Wang, Zhengwu Liu, Ngai Wong Wang Kang</i>	
Diagnostic Test Generation for Fault Localization in Printed Neuromorphic Circuits	793
<i>Tara Gheshlaghi, Alexander Studt, Priyanjana Pal, Dina A. Moussa, Michael Hefenbrock, Michael Beigl Mehdi B. Tahoori</i>	
GPU-Accelerated Efficient Transduction for Logic Optimization	796
<i>Zhuofan Lin, Shiju Lin</i>	

DynaOpt: A Heterogeneous Logic Optimization Framework with Dynamic Sequence Generation	803
<i>Xingyu Qin, Guande Dong, Jianwang Zhai, Kang Zhao</i>	
Exact Synthesis with Optimal Switching Activity	810
<i>Marcel Walter, Michael Feldmeier, Robert Wille</i>	
HoPart: Hop-Constrained Partitioning with Routing Support for Multi-FPGA Systems	817
<i>Yuan Huang, Longkun Guo, Weijie Fang, Jiawei Lin</i>	
QUADOL: A Quality-Driven Approximate Logic Synthesis Method Leveraging Dual-Output LUTs for Modern FPGAs	824
<i>Jian Shi, Chang Meng, Xuan Wang, Weikang Qian</i>	
MetaSyn: A Meta-Reinforcement Learning Framework with Multimodal Circuit Representation for Adaptive Logic Synthesis	831
<i>Shukai Liu, Ruoyan Liao, Siyu Wang, Qimin Xu, Cailian Chen</i>	
eLogic: An E-Graph-based Logic Rewriting Framework for Majority-Inverter Graphs	838
<i>Rongliang Fu, Wei Xuan, Shuo Yin, Guangyu Hu, Chen Chen, Hongce Zhang, Bei Yu, Tsung-Yi Ho</i>	
Advancing LUT-based Threshold Logic Synthesis with Enhanced Area Estimation	844
<i>Yu-Shan Lin, Yung-Chih Chen</i>	
UniCircuit: Multimodal Circuit Representation Learning with Anchor-Free Alignment	851
<i>Jingxin Wang, Weikang Qian</i>	
Efficient Arithmetic on FPGA	854
<i>Danila Gorodecky, Leonel Sousa</i>	
A ² RT: Efficient Ray Tracing Accelerator with Approximate-Accurate Computing and Quantization	856
<i>Zhiyuan Zhang, Zhihua Fan, Wenming Li, Yudong Mu, Yuhang Qiu, Zhen Wang, Xiaochun Ye, Xuejun An</i>	
KAN-SAs: Efficient Acceleration of Kolmogorov-Arnold Networks on Systolic Arrays	863
<i>Sohaib Errabii, Olivier Sentieys, Marcello Traiola</i>	
Scrooge: Accelerating Attention Inference in LLMs via Early Termination Mechanism	870
<i>Gwangeun Byeon, Seongwook Kim, Taein Kim, Jungmin Lee, Seokin Hong</i>	
FLICKER: A Fine-Grained Contribution-Aware Accelerator for Real-Time 3D Gaussian Splatting	877
<i>Wenhui Ou, Zhuoyu Wu, Yipu Zhang, Dongjun Wu, Freddy Ziyang Hong, C. Patrick Yue</i>	
RAPID: Accelerating Point Cloud Diffusion Models via Space-Aware Mix-Precision Quantization	884
<i>Qichu Sun, Yanan Zhu, Linxi Lu, Haishuang Fan, Jingya Wu, Huawei Li, Xiaowei Li, Guihai Yan</i>	
AO-BFP: An Adaptive Mixed-Precision and Outlier-Aware Block Floating-Point Accelerator for Large Language Model Inference	891
<i>Yifan Wang, Zetao Guo, Wendi Sun, Wenhao Sun, Qiyang Fang, Song Chen, Yi Kang</i>	
Breaking the BRAM Wall: Scalable VINA FPGA Acceleration via Distributed Grid Storage and Cross-Board Long-Ring Pipelines	898
<i>Ankun Tian, Shidi Tang, Ruiqi Chen, Ming Ling</i>	
AURORA - AUTomated 8T SRAM Wired-OR Logic Array for Boolean-Based Machine Learning	905
<i>Komal Krishnamurthy, Marcos L. L. Sartori, Shengyu Duan, Alex Yakovlev, Rishad Shafik</i>	

Microscaling–Stochastic Computing Based Systolic Arrays for Energy-Efficient Deep Neural Network Inference	912
<i>Mohammad Hassani Sadi, Bilal Hammoud, Norbert Wehn</i>	
CoMix-D: A Low-Cost, RNG-Free Decorrelator via Correlation Mixing for Stochastic Computing	915
<i>Yexian Lin, Chunyan Wu, Kuncai Zhong, Weikang Qian</i>	
Non-volatile Spintronic Flip-Flops with Checkpoint Preservation Supported in RISC-V Platform	918
<i>Jiongzhe Su, Mingtao Chen, Zhanpeng Qiu, Bo Liu, Hao Cai</i>	
A Low-Power Bayesian Head Using SOT-MRAM Arrays for Uncertainty-Aware Binary Neural Networks	925
<i>Joao Henrique Quintino Palhares, Bruno Lovison Franco, Louis Hutin, Jonathan Miquel Kamel-Eddine Harabi, Aymen Romdhane, Kevin Garelo</i>	
A Low Power and High Reliability Nonvolatile SRAM Using In-Plane VGSOT-MRAM with Pre-Charge Restore Scheme	932
<i>Xiaoyang Xu, Chenyi Wang, Zhongzhen Tong, Mingche Li, Weimeng Zhao, Zhongkui Zhang, Yaling Wang Chao Wang, Zhaohao Wang</i>	
Equivalent-Ons-Replacement Self-Aware-Access LLC on Dual-Port SOT-MRAM by Sense-While-Replace	939
<i>Keyang Zhang, Quanhai Zhu, Zhenghan Fang, Shuyu Wang, Hao Cai</i>	
Soft-Error Resilient MRAM-OTP BCAM for DDR4 STT-MRAM Redundancy Management	946
<i>Haoran Du, Hongjin Zhu, Zhenghan Fang, Shuyu Wang, Hao Cai</i>	
MIRAGE:MRAM-Based Near ADC-Less Compute-In-Memory Macro for Deep Learning Acceleration	953
<i>Mainakh Mukherjee, Ayan B. Pranta, Utkarsh Saxena, Anushka Mukherjee, Deepika Sharma Gaurav Kumar K, Kaushik Roy</i>	
An Effective SNN Macro with Real-Time STDP and Dynamic LIF Model Based on Thermally Interplayed Spin-Orbit Torque MTJ	960
<i>Changyu Li, Linjun Jiang, Liangchen Li, Dehang Zhu, Junda Zhao, Hongxi Liu, Wang Kang, Wenlong Cai He Zhang, Weisheng Zhao</i>	
Antiferromagnetic Tunnel Junctions (AFMTJs) for In-Memory Computing: Modeling and Case Study	967
<i>Yousuf Choudhary, Tosiron Adegbiya</i>	
High-Performance and High-Density NAND-Like SOT-MRAM for FinFET Technology Nodes	970
<i>Chao Wang, Xianzeng Guo, Luman Xiang, Zhaohao Wang, Weisheng Zhao</i>	
SecIC3: Customizing IC3 for Hardware Security Verification	973
<i>Qinhan Tan, Akash Gaonkar, Yu-Wei Fan, Aarti Gupta, Sharad Malik</i>	
INSPIRE: In-Sensor Compressed Weight Retrieval for Enhancing ViT Efficiency at Edge	980
<i>Sabbir Ahmed, Deniz Najafi, Mohaiminul Al Nahian, Navid Khoshavi, Abdullah Al Arafat Mamshad Nayeem Rizve, Mahdi Nikdast, Adnan Siraj Rakin, Shaahin Angizi</i>	
LoRA-Edge: Tensor-Train–Assisted LoRA for Practical CNN Fine-Tuning on Edge Devices	987
<i>Hyunseok Kwak, Kyeongwon Lee, Jae-Jin Lee, Woojoo Lee</i>	
Focus Session: Advanced CMOS and 5.5D Packaging: Perspectives and Challenges for Design, Reliability and Security	994
<i>Hussam Amrouch, Dragomir Milojevic, Giorgio Di Natale, Jérôme Toubanc</i>	

Special Day—A Design Blueprint for Scalable Multi-Agent Architectures in Complex EDA Workflows	1000
<i>Valerio Tenace, Pierre-Emmanuel Gaillardon</i>	
Special Day – GUIDE: GenAI Units In Digital Design Education	1007
<i>Weihua Xiao, Jason Blocklove, Matthew DeLorenzo, Johann Knechtel, Ozgur Sinanoglu, Kanad Basu</i> <i>Jeyavijayan Rajendran, Siddharth Garg, Ramesh Karri</i>	
When Faults Don't Vanish: Persistent Fault Injection and Key Recovery on MRAM-Backed AES	1014
<i>Brojogopal Sapui, Priyanjana Pal, Mehdi B. Tahoori</i>	
CAMI: A Context-Aware Isolation Architecture for GPU Memories	1021
<i>Hao Lan, Wei Yan, Qinfen Hao, Xiaochun Ye, Yier Jin, Yong Liu, Ninghui Sun</i>	
ProCamo: A Fast Post-Manufacturing Programmable Camouflaged Logic Family	1028
<i>Seo Hyun Kim, Minhyeok Jeong, Jongmin Lee</i>	
NuRedact: Non-Uniform eFPGA Architecture for Low-Overhead and Secure IP Redaction	1035
<i>Voktho Das, Kimia Azar, Hadi Kamali</i>	
Automated Hardware Trojan Insertion in Industrial-Scale Designs	1042
<i>Yaroslav Popryho, Debjit Pal, Inna Partin-Vaisband</i>	
On Oracle-Guided Random Circuit Learning via Stochastic Boolean Satisfiability	1049
<i>Shakil Ahmed, Kaveh Shamsi</i>	
COVERT: Trojan Detection in COTS Hardware via Statistical Activation of Microarchitectural Events	1056
<i>Mahmudul Hasan, Sudipta Paria, Swarup Bhunia, Tamzidul Hoque</i>	
Circuit-Aware Analysis of Arithmetic Error Detection Codes	1063
<i>Cheng Chiu, Keyon Mazandarani, Nathan Bleier</i>	
Provably Optimal Planar Pareto Nearest Neighbor Search with Double Monotone Chains	1066
<i>Zizheng Guo, Runsheng Wang, Yibo Lin</i>	
Near-Optimal TDM Ratio Assignment for Die-Level Routing in Multi-FPGA Systems	1073
<i>Jiawei Lin, Longkun Guo, Weijie Fang</i>	
DRVision: A DRV-aware Routability Optimization Framework with Multi-modal Prediction and Vision-Based Routing Guidance	1080
<i>Xu Cheng, Pengcheng Fan, Peng Cao</i>	
TANGRAM: A Novel ILP-based On-Track Bus Routing via Placement and Compression of Polygons	1087
<i>Jaekyung Im, Seokhyeong Kang</i>	
DSR: A Systematic Approach for Efficient Double-sided Signal Routing	1093
<i>Jianqing Chen, Zhenkun Lin, Xun Jiang, Genggeng Liu, Yibo Lin, Gang Du</i>	
Submodular Maximization-inspired Adaptive Routing Bend Space Planning	1100
<i>Siting Liu, Peng Xu, Peiyu Liao, Keren Zhu, Yibo Lin, Bei Yu</i>	
An Adaptive Cost-based Via and Congestion Co-optimization Framework for VLSI Global Routing	1107
<i>Zhaoyi Wu, Haishan Huang, Jianli Chen, Zhifeng Lin</i>	
Optimizing Multibit Flip-Flop Banking via Agile In-Placement PPA Co-Optimization	1114
<i>Huan-Yuan Chen, Yu-Ruei Lin, Mark Po-Hung Lin, Hung-Ming Chen</i>	

BADGE: Boundary-Aware Dirichlet Graph Embedding for Initialization in DREAMPlace	1117
<i>Jaemin Park, Taejin Paik</i>	
MX-SAFE: Versatile Inference-and Training-Proof Microscaling Format with On-the-Fly Exponent and Mantissa Bit Allocation	1120
<i>Dahoon Park, Jahyun Koo, Sangwoo Hwang, Jaeha Kung</i>	
PatBiNN: A 65 nm Processing-in-CAM Based BNN Implementation for Pathogen Genome Classification	1127
<i>Yuval Harary, Almog Sharoni, Esteban Garzón, Leonid Yavits</i>	
HI-APP: Hardware-friendly Fully-Integer Approximation of Nonlinear Functions in Quantized CLIP-ViTs	1134
<i>Beom Jin Kang, Hyun Kim</i>	
BOLD-Q: Blockwise Outlier-aware Logarithmic Dual-Bias Quantization for Hardware-Efficient LLM Inference	1141
<i>Sungsoo Han, Dahun Choi, Hyun Kim</i>	
T-SAR: A Full-Stack Co-design for CPU-Only Ternary LLM Inference via In-Place SIMD ALU Reorganization	1148
<i>Hyunwoo Oh, KyungIn Nam, Rajat Bhattacharjya, Hanning Chen, Tamoghno Das, Sanggeon Yun, Suyeon Jang, Andrew Ding, Nikil Dutt, Mohsen Imani</i>	
IterQuant: Iterative Quantization Framework for Mixed-Precision LLM Compression	1155
<i>Hyungyo Jeong, Jiwoo Kim, Hyeokjun Kwon, Jaeho Lee, Youngjoo Lee</i>	
GCPT: Gradient-aware Clustering Method for Efficient Post-Training Quantization in Large Neural Networks	1162
<i>Chuyi Dai, Chen Ye, Zeyu Li, Jun Tao, Wei Zhang, Grace Li Zhang, Xin Li</i>	
Fast and Energy-Efficient Support for Low-Precision LLMs on PIM	1169
<i>Byeori Kim, Sangjun Lee, Eunhyeok Park</i>	
HAP: Accelerating DNNs with Resolution-Preserved Quantization by Harnessing Adaptive-Precision	1176
<i>Erjing Luo, Xinkuang Geng, Honglan Jiang, Leibo Liu, Jie Han</i>	
ArchE-Q: A DSP-Free Dataflow Accelerator for Quantized Neural Networks in Sensor-Aided Millimeter-Wave Edge Connectivity	1179
<i>Arish Sateesan, Ljiljana Simić, Marina Petrova</i>	
Late Breaking Results: Practical Power Side-Channel Attack on Analog Compute-in-Memory Macro	1182
<i>Simon Wilhelmstätter, Johannes Stark, Devanshi Upadhyaya, Maël Gay, Ilia Polian, Maurits Ortmanns</i>	
Late Breaking Results: Boosting Efficient Dual-Issue Execution on Lightweight RISC-V Cores	1185
<i>Luca Colagrande, Luca Benini</i>	
Late Breaking Results: A Power-Efficient RISC-V Baseband System-on-Chip for Multi-Standard Integrated Sensing and Communications	1188
<i>Limin Jiang, Yi Shi, Yihao Shen, Yintao Liu, Siyi Xu, Qingyu Deng, Xiaoxiao Chen, Qianli Wang, Shan Cao, Zhiyuan Jiang, Sheng Zhou</i>	
Late Breaking Results: Float Fight - Verifying Floating-Point Behavior in RISC-V Simulators	1191
<i>Katharina Ruep, Manfred Schlägl, Daniel Große</i>	

Late Breaking Results: Efficient Formal Verification of Highly Optimized MAC Units	1194
<i>Jan Kleinekathöfer, Lennart Weingarten, Kamalika Datta, Rolf Drechsler</i>	
Late Breaking Results: PolyRAD - Polynomial Formal Verification of Restoring Array Dividers	1197
<i>Mohamed Nadeem, Chandan Kumar Jha, Rolf Drechsler</i>	
Late Breaking Results: CHESSY: Coupled Hybrid Emulation with SystemC-FPGA Synchronization	1200
<i>Lorenzo Ruotolo, Giovanni Pollo, Mohamed Amine Hamdi, Matteo Risso, Yukai Chen, Enrico Macii</i> <i>Massimo Poncino, Sara Vinco, Alessio Burrello, Daniele Jahier Pagliari</i>	
Late Breaking Results: Thermally Assisted RowPress-RowHammer Synergy for Cross-Row Bit Flips	1203
<i>Filip Roth Tronnes-Christensen, Ranyang Zhou, Gamana Aragonda, Abeer Matar A. Almalky</i> <i>Mohaiminul Al Nahian, Adnan Siraj Rakin, Shaahin Angizi</i>	
Late Breaking Results: Never-Stopping Inference: Self-Healing AI Accelerators on SRAM-FPGAs	1206
<i>Eleonora Vacca, Giorgio Cora, Luca Sterpone</i>	
Late Breaking Results: RL-Based Macro Placement with Cell Clustering and Rudy Modeling for Routability Optimization	1209
<i>Youwen Wang, Hao Gu, Xinglin Zheng, Keyu Peng, Ziran Zhu</i>	
Lithography Hotspot Detection for Complex Non-Manhattan Layouts via Graph Neural Network	1212
<i>Bohao Li, Ranran Liu, Yumeng Liu, Cong Jiang, Kang Liu, Bei Yu, Kun Ren, Qi Sun, Cheng Zhuo</i>	
RIFT: A Scalable Methodology for LLM Accelerator Fault Assessment using Reinforcement Learning	1218
<i>Khurram Khalil, Muhammad Mahad Khaliq, Khaza Anuarul Hoque</i>	
Smart-PCLib: A LLM-based Multi-Agent Framework for Automated PCB Component Library Generation	1225
<i>Zhaohai Di, Jindong Tu, Zhiyuan He, Yuan Pu, Jiawei Liu, Chong Tong, Tsung-Yi Ho, Bei Yu</i> <i>Tinghuan Chen</i>	
Fault-tolerance Mapping of Spiking Neural Networks to RRAM-based Neuromorphic Hardware	1232
<i>Yuqing Xiong, Cao Xiao, Zhijie Yang, Lei Wang, Mengying Zhao</i>	
MeshHD: Near-Linear Encoding for Hyperdimensional Computing via Multi-Scale Bases and Kronecker Factorization	1239
<i>Woong Jae Han, Jiseung Kim, Hyukjun Kwon, Hojeong Kim, Selim An, Shinhyoung Jang, Yeseong Kim</i>	
Learning to Sense Without ADCs: Exploiting Phasic Responses from Diffusive Memristors	1246
<i>Zhenhang Zhang, Jingang Jin, Ruoyu Zhao, Zixu Wang, Tong Wang, Rui Zuo, J. Joshua Yang, Qinru Qiu</i>	
TSIM4ICS: Trace-Driven SystemC-TLM Simulation Framework for I/O Die-Based Multi-Chiplet Systems	1253
<i>Youngchul Yoon, Soonhoi Ha</i>	
FlashGEMM: Mesh-Aware Efficient GEMM for 3D-Stacked LLM Accelerators	1260
<i>Xin Fan, Chen Bai, Xin Yang, Zhenhua Zhu, Yanhong Wang, Zhaode Zhang, Yuan Xie</i>	
X-Sim: An Accurate and Scalable Simulator for Memristive Computing-in-Memory Accelerators	1267
<i>Konstantinos Stavarakakis, Bas Smeele, Emmanouil Arapidis, Theofilos Spyrou, Anteneh Gebregiorgis</i> <i>Stephan Wong, Georgi Gaydadjiev, Said Hamdioui</i>	
Graph-SRAM: Efficient Graph Learning-based SRAM Simulation via Waveform Propagation	1274
<i>Beisi Lu, Lihao Liu, Li Shang, Fan Yang</i>	

MGPA: A Memristor-based Genome Processing Accelerator for Single-cell RNA Sequencing	1281
<i>Yang Han, Lianfeng Yu, Teng Zhang, Bowen Wang, Yihang Zhu, Lei Cai, Yaoyu Tao, Yuchao Yang</i>	
Synthesizing Mixed-Mode Operations for Memristors using Majority Decomposition	1284
<i>Felix Bayhurst, Li-Wei Chen, Kefeng Li, Heidemarie Krüger, Nan Du, Ilia Polian</i>	
MinFill: Reinforcement Learning and GNN Guided Reordering for Fill-In Reduction in RF Circuit Matrices	1287
<i>Hao Zhang, Dan Niu, Cheng Zhuo, Zhou Jin</i>	
IP-Matcher: An Efficient One-to-Many Matching Framework for Analog Circuit Design and Reusing	1294
<i>Shixin Chen, Peng Xu, Yapeng Li, Tinghuan Chen, Bei Yu</i>	
RASNIL: PVT-Robust Many-Objective Analog Sizing via Nested Hybrid Fidelity Framework with Incremental Learning	1301
<i>Xingyu Tang, Sen Yin, Zhujun Yao, Bingzhang Huang, Xiaosen Liu, Yan Wang</i>	
Substrate: A Statically Typed Framework for Designing Highly Configurable Analog and Mixed-Signal Circuit Generators	1308
<i>Rahul Kumar, Rohan Kumar, Borivoje Nikolić</i>	
ACEMARL: Adaptive Clustering Enhanced Multi-Agent Reinforcement Learning for Analog Circuit Sizing	1315
<i>Han Wu, Haoning Jiang, Zhuoli Ouyang, Ziheng Wang, Qi Shen, Bo Yuan, Yan Lu, Junmin Jiang</i>	
Simulator-Driven Deep Reinforcement Learning for Analog Circuit Design	1322
<i>Felicia B. Guo, Ken T. Ho, Andrei Vladimirescu, Borivoje Nikolić</i>	
AutoPMS: A Framework for Automated Power Management System Design via Hierarchical Modeling and Embedded Design Knowledge	1329
<i>Bin Ye, Shuo Li</i>	
CaDRO: Causal-Guided Dimensionality Reduction for Scalable Multi-Objective Pareto Optimization	1336
<i>Dinithi Jayasuriya, Divake Kumar, Sureshkumar Senthilkumar, Devashri Naik, Nastaran Darabi Amit Ranjan Trivedi</i>	
DiffResist: Physics-Constrained Diffusion for Photoresist Modeling	1343
<i>Zixiao Wang, Jieya Zhou, Xinyun Zhang, Shoubo Hu, Farzan Farnia, Bei Yu</i>	
GRAIN: A Design-Intent-Driven Analog Layout Migration Framework	1346
<i>Bingyang Liu, Haoning Jiang, Haoyi Zhang, Xiaohan Gao, Zichen Kong, Xiyuan Tang, David Z. Pan Yibo Lin</i>	
LATIAS: A General Architecture-Operator Model for Spatial Accelerators with Complex Topology and Memory Hierarchy	1349
<i>Chengrui Zhang, Liancheng Jia, Chu Wang, Tianqi Li, Renze Chen, Xiuping Cui, Size Zheng, Shengen Yan Xiuhong Li, Yu Wang, Xiang Chen, Yun Eric Liang</i>	
CHIP-MAP: A Collaborative Optimization Framework for Macro Placement Using Large Language Models	1356
<i>Yiming Du, RenYe Yan, Yunfan Yang, Frank Qu, Jiajun Tan, ZhiYu Zheng, Yiming Gan, Ling Liang Zongwei Wang, YiMao Cai</i>	
LithoMamba: High-fidelity lithography simulation with State Space Models	1363
<i>Xinyu He, Daohui Wang, Shujing Lyu, Pourya Shamsolmoali, Jiwei Shen, Yue Lu</i>	

Explainable Hardware Trojan Detection at RTL using Attention Mechanism	1369
<i>Siyu Tian, Wei Hu, Lingjuan Wu, Tianle You, Hao Su, Jiacheng Zhu</i>	
MAEDA: An LLM-Powered Multi-Agent Evaluation Framework for EDA Tool Documentation QA	1376
<i>Zhenghao Chen, Yuan Pu, Hairuo Han, Yuntao Nie, Jiajun Qin, Yuhao Qin, Tairu Qiu, Zhuolun He</i> <i>Jianwang Zhai, Bei Yu, Kang Zhao</i>	
ChatTest: Coverage-Enhanced Testbench Generation for Agile Hardware Verification with LLMs	1383
<i>Gwok-Waa Wan, Shengchu Su, Jingyi Zhang, Sam Zaak Wong, Mengnv Xing, Lei Ji, Zhe Jiang, Xi Wang</i> <i>Jun Yang</i>	
GANGR: GAN-Assisted Scalable and Efficient Global Routing Parallelization	1390
<i>Hadi Khodaei Jooshin, Inna Partin-Vaisband</i>	
EstCoder: A RTL Code Generator based on Static Functional Estimation	1397
<i>Qi Xiong, Renzhi Chen, Zhigang Fang, Bowei Wang, Yingjie Zhou, Libo Huang, Lei Wang</i>	
FAST: A Scalable Framework for Accelerating Flexible Structured Sparse Training	1400
<i>Shuaiheng Li, Jun Liu, Xinhao Li, Yaoxiu Lian, Tianlang Zhao, Li Ding, Guohao Dai</i>	
DARE: An Irregularity-Tolerant Matrix Processing Unit with a Densifying ISA and Filtered Runahead Execution	1402
<i>Xin Yang, Xin Fan, Zengshi Wang, Jun Han</i>	
XTree on EquiMesh: Topology and Algorithm Co-Design for Collective Communication	1409
<i>Junwei Cui, Le Qin, Weilin Cai, Jiayi Huang</i>	
Beaivi: A 22-nm 1-GHz+ Exposed Datapath RISC-V DSP for Low-Power Applications	1416
<i>Kari Hepola, Joonas Multanen, Väinö-Waltteri Granat, Jakub Žádník, Roope Keskinen, Karri Palovuori</i> <i>Pekka Jääskeläinen</i>	
T-MSA: Transformer-Driven Multi-Strategy Adaptive Microarchitecture Design Space Exploration	1423
<i>Jingjing Wang, Zihan Lin, Fan Yang, Xiaochuan Li, Runze Zhang, Cong Xu, Rengang Li, Baoyu Fan</i>	
RISC-V ISA Extensions for Vectorized Unstructured Sparse SpMM in LLM Inference	1430
<i>Tengfei Xia, Zhihua Fan, Jing Xue, Shantian Qin, Xiaochun Ye, Wenming Li</i>	
From Buffers to Registers: Unlocking Fine-Grained FlashAttention with Hybrid-Bonded 3D NPU Co-Design	1437
<i>Jinxin Yu, Yudong Pan, Mengdi Wang, Huawei Li, Yinhe Han, Xiaowei Li, Ying Wang</i>	
Increasing the Efficiency of Associative Processor Architectures via CMOS-Compatible Hybridization	1444
<i>Socrates S. Wong, Cecilio C. Tamarit, Mohammad M. Sharifi, Zephan Enciso, Dayane Reis</i> <i>Michael Niemier, X. Sharon Hu, José F. Martínez</i>	
PQCUARK: A Scalar RISC-V ISA Extension for ML-KEM and ML-DSA	1451
<i>Xavier Carril, Alicia Manuel Pasoot, Emanuele Parisi, Oriol Farràs, Carlos Andrés Lara-Niño</i> <i>Miquel Moretó</i>	
NITRO: 3D NAND Flash-Based In-Storage LLM Computing with Enhanced Activation Dataflow	1458
<i>Sanghun Shin, Gisan Ji, Sungju Ryu</i>	
SMIX: Schedulable Instruction Set Architecture Extension Interface for Multi-Operand Operators	1461
<i>Shufan He, Hanmo Wang, Kefa Chen, Xuyin Chen, Xianhua Liu, Chun Yang</i>	

Multi-Partner Project: COIN-3D – Collaborative Innovation in 3D VLSI Reliability	1464
<i>George Rafael Gourdoumanis, Fotini Oikonomou, Maria Pantazi-Kypraiou, Pavlos Stoikos</i>	
<i>Olympia Axelou, Athanasios Tziouvaras, Georgios Karakostas, Tahani Aladwani</i>	
<i>Christos Anagnostopoulos, Yixian Shen, Anuj Pathania, Alberto Garcia-Ortiz, George Floros</i>	
Multi-Partner Project: Multi-GPU Performance Portability Analysis for CFD Simulations at Scale	1470
<i>Panagiotis-Eleftherios Eleftherakis, George Anagnostopoulos, Anastassis Kapetanakis, Mohammad Umair</i>	
<i>Jean-Yves Vet, Konstantinos Iliakis, Jonathan Vincent, Jing Gong, Akshay Patil, Clara García-Sánchez</i>	
<i>Gerardo Zampino, Ricardo Vinuesa, Sotirios Xydis</i>	
Multi-Partner Project: CeCaS Accelerator Design for Efficient Supercomputing in Automotive Systems	1476
<i>Annina Gutermann, Alexey Serdyuk, Fabian Lesniak, Julian Hoefer, Hella Toto Kies, Tanja Harbaum</i>	
<i>Juergen Becker, Brian Pachideh, Sven Nitzsche, Moritz Neher, Carmen Weigelt, Jann Krausse</i>	
<i>Victor Pazmino, Klaus Knobloch, Lukas Groth, Andrija Nešković, Saleh Mulhem, Mladen Berekovic</i>	
Multi-partner project: Ecodesign to reduce electronic waste	1483
<i>Chiara Sandionigi, Olivia Belorgeot, Elise Chumat, Jean-Christophe Crebier, George Dimitrakopoulos</i>	
<i>Lena Froger, Thomas Krivec, Michel Monsellier, Lucas Pinto, Ioannis Vondikakis, Christof Wernbacher</i>	
<i>Edmund Whitmore, Olivier Pedoussaut</i>	
Multi-Partner Project: Scheduling-Deployment Workflow for Autonomous RoboRacer Driving Stacks in the HAL4SDV Project	1490
<i>Matthias Stammler, Henrik Scheidt, Tanja Harbaum, Jürgen Becker, Konstantin Dudzik, Victor Pazmino</i>	
<i>Federico Gavioli, Paolo Burgio, Arvind Easwaran, Andreas Eckel</i>	
Multi-Partner Project: Advancing European Semiconductor and Chiplet Innovation Through the Bavarian Chip Design Center	1497
<i>Hussam Amrouch, Jehaan Joseph, Michael Schirmer, Johannes Geier, Ulf Schlichtmann</i>	
<i>Michael Meidinger, Thomas Wild, Andreas Herkersdorf, Jens Nöpel, Georg Sigl, Yicheng Zhang</i>	
<i>Carsten Trinitis, Aswathy Nedumpalli Sankaranarayanan, Martin Schulz, Andreas Korb</i>	
<i>Konrad Hohentanner</i>	
Multi Partner Project: STRATUM, co-creation protocol and advanced smart GUI for a 3D neurosurgery supporting tool	1504
<i>Emanuele Torti, Himar Fabelo, Elisa Marenzi, Maria Luisa Alvarez-Male, Chrysanthi Bairaktari</i>	
<i>Beatriz Noriega-Ortega, Raquel Leon, Santiago Marco, Asaf Badouh, Max Verbers, Carlos Vega</i>	
<i>Javier Santana-Nuñez, Yolanda Ramallo-Fariña, Christian Weis, Ana M Wäagner, Eduardo Juarez</i>	
<i>Claudio Rial, Alfonso Lagares, Gustav Burström, Luis Jimenez-Roldan, Teresa Cervero, Miquel Moreto</i>	
<i>Giovanni Danese, Svitlana Zinger, Francesca Manni, Miguel A. García-Bello, Lidia García, Jesus Morera</i>	
<i>Juan F. Piñeiro, Bernardino Clavo, Francesco Leporati, Gustavo M. Callico</i>	
Multi-Partner Project: UP2DATE4SDV : Enabling Safe and Secure Modular Updates, Upgrades and Dynamic Task-Reallocation and -Execution for Software Defined Vehicles	1511
<i>Gregor Nitsche, Patrick Uven, Hannes Fuchs, Enrico Mezzetti, Marcus Hähnel, Mijangos Ane</i>	
<i>Kim Grüttner</i>	
Focus Session: Advanced Hybrid Hardware Fuzzing	1518
<i>Chen Chen, Stephen Muttathil, Mohamadreza Rostami, Nikhilesh Singh, Lichao Wu, Ahmad-Reza Sadeghi</i>	
<i>Jeyavijayan Rajendran</i>	
Focus Session: What the Fuzz! Pushing Beyond Randomness in Hardware Security with Generative AI	1525
<i>Nikhilesh Singh, Mohamadreza Rostami, Lichao Wu, Chen Chen, Stephen Muttathil, Jeyavijayan Rajendran</i>	
<i>Ahmad-Reza Sadeghi</i>	

Focus Session: Exploring Semantic Leakage in Edge FPGA Implementations of Neural Networks	1532
<i>Zhuoran Liu, Konstantina Miteloudi, Durba Chatterjee, Lejla Batina</i>	
RLConcolic: Enhancing Concolic Testing via Multi-Step Reinforcement Learning	1539
<i>Yan Tan, Xiangchen Meng, Yangdi Lyu</i>	
Fine-Grained Code Analysis for Processor Fuzzing	1546
<i>Ziyue Zheng, Zhi Qu, Yangdi Lyu</i>	
Validating Formal Hardware Specifications Through Generated Behavioral Models	1553
<i>Robert Kunzelmann, Zeyad Tahoun, Vinod Bangalore Ganesh, Maximilian Berger, Emil Baerens Wolfgang Ecker</i>	
FORWORD: Accelerating Formal Datapath Verification via Word-Level Sweeping	1560
<i>Ziyi Yang, Guangyu Hu, Xiaofeng Zhou, Mingkai Miao, Changyuan Yu, Wei Zhang, Hongce Zhang</i>	
VLIM: Verified Loop Interchange for Optimised Matrix Multiplication	1567
<i>Oliver Turner, Shounak Chakraborty</i>	
Polynomial Verification of 2-Affine Spaces	1574
<i>Anna Bernasconi, Valentina Ciriani, Gianmarco Cuciniello, Caroline Dominik, Rolf Drechsler</i>	
Hyperplane Input Space Cuts for Neural Network Verification	1581
<i>Jonathan Hjort, Ahmed Rezone</i>	
Provable Guarantees in Approximate Synthesis	1588
<i>Kushagra Gupta, Priyanka Golia, Subhajit Roy, Kuldeep S Meel</i>	
A Formally Verified Secure Caching Mechanism on TrustZone-enabled Microcontrollers	1595
<i>Salvatore Bramante, Matteo Busi, Alessandro Cilardo, Riccardo Focardi, Flaminia Luccio Stefano Mercogliano</i>	
DecoHD: Decomposed Hyperdimensional Classification under Extreme Memory Budgets	1598
<i>Sanggeon Yun, Hyunwoo Oh, Ryoze Masukawa, Mohsen Imani</i>	
FaTRQ: Tiered Residual Quantization for LLM Vector Search in Far-Memory-Aware ANNS Systems	1605
<i>Tianqi Zhang, Flavio Ponzina, Tajana Rosing</i>	
Entropy Sampling-Based Neural Architecture Search for Resource-Constrained Microcontroller Targets	1612
<i>Christian Heidorn, Frank Hannig, Dominik Riedelbauch, Christoph Strohmeier, Jürgen Teich</i>	
Toward Parallel Serving for Vision-Language Models via Modal Decoupling and Scheduling	1619
<i>Yijia Yang, Yubo Deng, Yida Wang, Yuanchao Xu, Keni Qiu</i>	
Dolphium: Co-Optimizing Quantization Dataflow and Paradigms on Poly-Hierarchical NPUs	1626
<i>Xiuping Cui, Chengrui Zhang, Xiang Chen, Yun Eric Liang</i>	
DynaMo: Runtime Switchable Quantization for MoE with Cross-Dataset Adaptation	1633
<i>Zihao Zheng, Xiuping Cui, Size Zheng, Maoliang Li, Jiayu Chen, Yun Liang, Xiang Chen</i>	
From Cloud-Heavy to Edge-Ready: Self-supervised Transfer-efficient Emotion Recognition	1640
<i>Junjiao Sun, José Miranda, Jorge Portilla, Andrés Otero</i>	
High-Efficiency Neural Beamforming for Real-Time Speech Enhancement on Smart Low-Power Hearable Devices	1643
<i>Luca Bompani, Giovanni Oltrecolli, Marco Fariselli, Francesco Conti</i>	

PCB-Migrator: Automated PCB PnR Migration	1646
<i>Yaohui Han, Beichen Li, Rongliang Fu, Qunsong Ye, Zhiyuan Lu, Junchen Liu, Bei Yu, Tsung-Yi Ho Tinghuan Chen</i>	
Dynamic Algorithm Configuration for Global Placement	1653
<i>Chen Lu, Ke Xue, Ruo-Tong Chen, Yunqi Shi, Siyuan Xu, Mingxuan Yuan, Chao Qian, Zhi-Hua Zhou</i>	
Timing-Driven Detailed Placement with Collaborative Topology Reconstruction	1660
<i>Zhengjie Zhao, Wenxin Yu, Jie Ma, Mengshi Gong, Youzhi Zheng, Xinmiao Li, Wenyu Liu, Jingwei Lu</i>	
HGNN-Part: A High-Quality Hypergraph Partitioner Based on Hypergraph Generative Model	1667
<i>Shengbo Tong, Rufan Zhou, Chunyan Pei, Wenjian Yu</i>	
HPPlacer: A High-Precision Slack-Aware Global Placement Engine	1674
<i>Qinggong Shen, Chaoli Zhang, Haoyang Xu, Zhiwen Yu, Bin Guo, Yuxuan Zhao, Bei Yu, Tsung-Yi Ho Xing Huang</i>	
Timing-driven Detailed Placement via TimingMask-guided Path-level Optimization	1681
<i>Ruo-Tong Chen, Chengrui Gao, Siyuan Xu, Ke Xue, Yunqi Shi, Xi Lin, Mingxuan Yuan, Chao Qian Zhi-Hua Zhou</i>	
3D Chiplet Partitioning and Floorplanning Interaction with Vertical Bonding Consideration	1688
<i>Tong Shen, Mengchen Chen, Xu He, Yao Wang, Yang Guo</i>	
PPA-driven Placement via Adaptive Cluster Constraints Optimization	1695
<i>Ziyan Liu, Siyuan Xu, Jie Wang, Zijie Geng, Yeqiu Chen, Mingxuan Yuan, Jianye Hao, Feng Wu</i>	
Towards learning-based gate-level glitch analysis	1702
<i>Anastasis Vagenas, Dimitrios Garyfallou, George Stamoulis</i>	
Drafting and Multi-Input Switching in Digital Dynamic Timing Simulation for Multi-Input Gates	1705
<i>Arman Ferdowsi, Ulrich Schmid, Josef Salzmann</i>	
Node2Node: Node Adaptation with Transformer for Cross-Node Hotspot Detection	1708
<i>Wenbo Xu, Silin Chen, Yibo Huang, Xinyun Zhang, Zixiao Wang, Bei Yu, Ningmu Zou</i>	
Population coding to improve fault tolerance of neuromorphic networks in regression tasks	1715
<i>Alexis Gleyo, Bernard Girau</i>	
Scalable Second-Order Optimizer for Full-Chip Inverse Lithography Techniques	1722
<i>Su Zheng, Ziyang Yu, Bei Yu, Martin Wong</i>	
Ramen: Radiation-Aware Modeling Framework for PDK-Enabled Design and Library Characterization	1729
<i>Quan Cheng, Haoyuan Li, Zhenzhe Chen, Wang Liao, Jing-Jia Liou, Masanori Hashimoto, Longyang Lin</i>	
No Pixel, More Efficient: Multimodal Framework for Sub-nm Mask Process Correction	1736
<i>Kai Ma, Tianyi Li, Jiaqi Liu, Jingyi Yu, Hao Geng</i>	
TRACE: A Transferable Framework for Aging-aware Cell Delay Estimation	1743
<i>Muyan Jin, Chao Yang, Yunlin Liu, Zejian Cai, Pengpeng Ren, Zhigang Ji</i>	
An IR drop-robust Mapping Method for Reliable Memristive Accelerators	1750
<i>Jinpeng Liu, Shiyi Song, Bing Wu, Huan Cheng, Heng Zhou, Xueliang Wei, Wei Tong, Dan Feng</i>	
Interpretable Graph Neural Networks for Fault Detection in Circuit Netlists	1757
<i>Rupesh Raj Karn, Johann Knechtel, Ozgur Sinanoglu</i>	

Analysis and Mitigation of IR Drop in Memristor-based AI Hardware Accelerators	1760
<i>Emmanouil Arapidis, Theofilos Spyrou, Konstantinos Stavrakakis, Emmanouil Anastasios Serlis</i>	
<i>Moritz Fieback, Said Hamdioui, Anteneh Gebregiorgis</i>	
An Efficient Weight Correction Method to Recover Non-ideal Errors in Pruned IRC Designs	1763
<i>Shih-Han Chang, Yi-Min Pan, Chong-En Hong, Chien-Nan Jimmy Liu</i>	
A Graph-Theoretic Framework for Randomness Optimization in First-Order Masked Circuits	1766
<i>Dilip Kumar S. V, Benedikt Gierlichs, Ingrid Verbauwhede</i>	
Extending and Accelerating Inner Product Masking with Fault Detection via Instruction Set Extension	1773
<i>Songqiao Cui, Geng Luo, Junhan Bao, Josep Balasch, Ingrid Verbauwhede</i>	
Rejection Matters: Efficient Non-Profiling Side-Channel Attack on ML-DSA via Exploiting Public Templates	1780
<i>Yuhan Zhao, Wei Cheng, Zehua Qiao, Yuejun Liu, Yongbin Zhou</i>	
Black-Box Robustness Probing of Graph Neural Networks for VLSI Circuit Netlists	1787
<i>Rupesh Raj Karn, Johann Knechtel, Ozgur Sinanoglu</i>	
Compact Yet Fast: An Efficient d-Order Masked Implementation of Ascon	1794
<i>Mattia Mirigaldi, Nico Paninfori, Maurizio Martina, Guido Masera</i>	
HEED: A Highly Efficient Electromagnetic Fault Detection Scheme	1801
<i>Roukoz Nabhan, Mohammad Ebrahimabadi, Jean-Luc Danger, Jean-Max Dutertre, Sylvain Guilley</i>	
<i>Naghme Karimi, Raphael Viera, Iyad Zaarour</i>	
GPU Acceleration of the Sum-Check Protocol Over Towers of Binary Fields for Verifiable Computing	1808
<i>Andrew Fan, Yanze Wu, Harry Han, Md Tanvir Arafin</i>	
Formal Methods-Assisted Chosen Ciphertext Attacks on PQC CRYSTALS-Kyber Using Electromagnetic Emanations	1815
<i>Yashaswini Makaram, Davis Ranney, Adam A. Ding, David Kaeli, Yungsi Fei</i>	
Synthesizable PUF Design with Library Characterization for Secure Storage in Edge Devices	1822
<i>Yuseong Lee, Donghyun Park, Jang Hyun Kim, Jongmin Lee</i>	
Optimization of AND-Gate-Sparse Circuit Synthesis for Multi-party Computation Systems with Local Communications	1825
<i>Yinfan Zhao, Makoto Ikeda</i>	
LiveVerilogEval: Contamination Free and Automatically Scalable Benchmark for Verilog Code Generation	1828
<i>Charles Young, Hao Yu, Dezhi Ran, Qingchen Zhai, Tianqi Qiu, Frank Qu, Bangyan Wang, Yuan Xie</i>	
<i>Tao Xie</i>	
ACES: A Chiplet Architecture with Resource Partition and Dynamic Scheduling for Agentic LLMs	1835
<i>Hongou Li, Mingxuan Li, Zhantong Zhu, Tianyu Jia</i>	
MeltRTL: Multi-Expert LLMs with Inference-time Intervention for RTL Code Generation	1842
<i>Nowfel Mashnoor, Mohammad Akyash, Hadi Kamali, Kimia Azar</i>	
MemoryIslands: A Federated Approach for Efficient Memory Mappings	1849
<i>Fatemeh Derakhshani, Mohamed Hassan</i>	

DPUConfig: Optimizing ML Inference in FPGAs Using Reinforcement Learning	1856
<i>Alexandros Patras, Spyros Lalis, Christos D. Antonopoulos, Nikolaos Bellas</i>	
A Collaborative Framework for Multi-Level Multi-Objective Design Space Exploration	1863
<i>Jiangnan Li, Kaixiang Zhu, Yuping Bai, Yunfei Dai, Qing He, Yu He, Lingli Wang</i>	
SINA: A Circuit Schematic Image-to-Netlist Generator Using Artificial Intelligence	1870
<i>Saoud Aldowaish, Yashwanth Karumanchi, Kai-Chen Chiang, Soroosh Noorzad, Morteza Fayazi</i>	
StreamNTT: A High-Throughput, HLS-Based Streaming NTT Accelerator for HBM-Equipped FPGAs	1873
<i>Young-Kyu Choi, Hyunwoo Park, Wei He, Sunwoong Kim</i>	
RouterAcc: FPGA Acceleration for VLSI Detailed Router via Hierarchical Storage Mapping	1880
<i>Ruiyuan Guo, Zexu Zhang, Chang Liu, Da Tang, Weiqi Shen, Haodong Lu, Xiqiong Bai, Kun Wang</i> <i>Jianli Chen, Jun Yu</i>	
Compacted-LUT: Fine-Grained Customizable LUT Architecture via SRAM-MUX Co-Optimization	1887
<i>Mingyang Chen, Yunfei Dai, Wai-Shing Luk, Qing He, Yu He, Lingli Wang</i>	
FLAME: A Framework Exploring Execution Strategies for Multi-Cycle Operations in CGRA	1894
<i>Jiajun Qin, Cheng Tan, Ruihong Yin, Tianhua Xia, Sai Qian Zhang, Bei Yu</i>	
NX-CGRA: A Programmable Hardware Accelerator for Core Transformer Algorithms on Edge Devices	1901
<i>Rohit Prasad</i>	
MIQARA: Mixed-Criticality Queue-based Architecture for Reconfigurable Accelerator Platforms	1908
<i>Hassan Nassar, Martin Rapp, Lars Bauer, Mostafa Elshimy, Zeynep Guelbeyaz Demirdag, Jörg Henkel</i>	
MC-CGRA: A Memory-Computation Coordinated CGRA Framework for Stream Processing	1915
<i>Chen Shi, Chunhua Xiao, Han Diao, Weijie Yuan, Junling Wang</i>	
CUT-MC: Optimizing the Relationship Between Context Count, Unrolling and Throughput in Multi-Context CGRAs	1921
<i>Stephen Wicklund, Jason Anderson</i>	
RIFT: A Single-Bitstream, Runtime-Adaptive FPGA-Based Accelerator for Multimodal AI	1924
<i>Hyunwoo Oh, Hanning Chen, Sanggeon Yun, Yang Ni, Suyeon Jang, Behnam Khaleghi, Fei Wen</i> <i>Mohsen Imani</i>	
A Cluster-Based Distributed Memory Architecture for CGRAs	1927
<i>Shangkun Li, Cheng Tan, Zeyu Li, Jinming Ge, Jiawei Liang, Hao Yang, Linfeng Du, Jiang Xu, Wei Zhang</i>	
Mixed-Precision Training and Compilation for RRAM-based Computing-in-Memory Accelerators	1930
<i>Rebecca Pelke, Joel Klein, José Cubero-Cascante, Nils Bosbach, Jan Moritz Joseph, Rainer Leupers</i>	
How to Manage the Mapping Table of Large-Capacity Solid-State Drives	1937
<i>Yang Zhou, Fang Wang, Dan Feng</i>	
FedCM: Fine-grained Kernel Scheduling and Management to Improve GPU Sharing	1944
<i>Xianwei Zhang, Xuanteng Huang, Nong Xiao</i>	
RheoSparse: Exploring Finer Grained Structured Sparsity for Small Language Models	1951
<i>Jianing Zheng, Gang Chen</i>	
Software-Based Approximate Multiplication on Multiplierless CPUs Using Custom Instruction	1958
<i>Shalu Prathmesh Rajiv, Rajesh Kedia</i>	

DyGen: A Constant-Time Kernel Generator for Dynamic-Shape Neural Networks	1965
<i>Yuhan Kang, Wenrui Zhang, Dong Chen, Yang Shi, Jianchao Yang, Zeyu Xue, Jing Feng, Mei Wen</i>	
SSR: Sparse Segment Reduction for Ternary GEMM Acceleration	1972
<i>Adeline Pittet, Shien Zhu, Valérie Verdan, Gustavo Alonso</i>	
EEMU: A FEMU-based Accurate, Parametric Ethernet-SSD Emulator	1979
<i>Xikun Jiang, Chao Li, Tianyu Wang, Xiaowei Chen, Zhaoyan Shen, Zili Shao</i>	
An Efficient Secure Boot Mechanism Leveraging DICE as a Use Case	1981
<i>Utku Budak, Malek Safieh, Yigit Arda Ozen, Fabrizio De Santis, Georg Sigl</i>	
Abusing DDS Discovery: Denial-of-Service Attacks Against ROS 2	1984
<i>Jiafu Xu, Songran Liu, Zilong Wang, Minghe Yu, Yue Tang, Yang Wang, Weiguang Pang, Wang Yi</i>	
Focus Session: Stepping Stones Towards Domain Acceleration Using AI-Driven High-Level Synthesis Design	1987
<i>Stefan Abi-Karam, Miaoyan Zhou, Cong "Callie" Hao</i>	
Focus Session: Do Agentic LLMs Change the Paradigm of Hardware Test Generation?	1993
<i>Farshad Firouzi, Pragya Sharma, Agastya Seth, Peter Domanski, Bahar Farahani, Sanmitra Banerjee</i>	
<i>Jonti Talukdar, Krishnendu Chakrabarty</i>	
Focus Session: Large Language Models in Physical Design: From Data Generation to Intelligent Agents	2000
<i>Bing-Yue Wu, Atmadip Dey, Austin Rovinski, Vidya A. Chhabria</i>	
Physical-Aware eFPGA Redaction for Secure and Efficient Hardware IP Protection	2007
<i>Yunqi He, You Li, Ruofan Huang, Guannan Zhao, Hai Zhou</i>	
The PMP Snapshot Engine: Fast and Fault-Resilient PMP Reconfiguration for RISC-V	2014
<i>Christian Larmann, Abdullah Aljuffri, Adrian Marotzke, Alejandro Garza, Said Hamdioui</i>	
<i>Mottaqiallah Taouil</i>	
A Reliability-Physics-Based Approach for Data Tampering Detection in Commercial 3D-NAND Flash Memory	2021
<i>Yuhan Wang, Jian Huang, Ruibin Zhou, Yao Liu, Haotian Ye, Xianping Liu</i>	
TrustSeed: Lightweight Attestation Protocol for Ensuring LLM Integrity	2028
<i>Mohamed Alsharkawy, Mohamed Aboelenien Ahmed, Hassan Nassar, Jeferson Gonzalez, Heba Khdr</i>	
<i>Osama Abboud, Xun Xiao, Jörg Henkel</i>	
Future Result Capture: Timing Anomalies Reveal Data from Instructions in the Successor	2035
<i>Roua Boulifa, Marwa Chehab, Paolo Maistri, Giorgio Di Natale</i>	
DEEP-LENS: Deep-Learning Powered Layout Extraction and Novel Segmentation for IC Assurance and Security	2042
<i>Shuvodip Maitra, Abhishek Chakraborty, Debdeep Mukhopadhyay</i>	
LAsset: An LLM-assisted Security Asset Identification Framework for System-on-Chip (SoC) Verification	2049
<i>Md Ajoad Hasan, Dipayan Saha, Khan Thamid Hasan, Nashmin Alam, Azim Uddin, Sujan Kumar Saha</i>	
<i>Mark Tehranipoor, Farimah Farahmandi</i>	

IMS: Intelligent Hardware Monitoring System for Secure SoCs	2056
<i>Wadid Foudhaili, Aykut Rencher, Anouar Nechi, Rainer Buchty, Mladen Berekovic, Andres Gomez Saleh Mulhem</i>	
Unlocking Hidden Secrets: Leveraging SRAM Aging Imprints for Sensitive Data Recovery	2063
<i>Zakia Tamanna Tisha, Gaines Odom, Biswajit Ray, Ujjwal Guin</i>	
Open-Source Framework for Secure Hardware Design with Simulation-based Leakage Assessment	2066
<i>Pablo Navarro-Torrero, Francisco J. Rubio-Barbero, Eros Camacho-Ruiz Macarena C. Martinez-Rodriguez, Piedad Brox-Jiménez</i>	
Design and Optimization of Mixed-Kernel Mixed-Signal SVMs for Flexible Electronics	2069
<i>Florentia Alentaki, Maha Shatta, Konstantinos Balaskas, Georgios Panagopoulos, Georgios Zervakis Mehdi B. Tahoori</i>	
A Multi-Sensor Approach for Soft Labeling in Human Activity Recognition Domain	2076
<i>Matteo Iervasi, Cristian Turetta, Florenc Demrozi, Graziano Pravadelli</i>	
LiteDVS: A Low-Data-Redundancy Dynamic Vision Sensor with Hybrid Readout and In-Pixel Denoising	2083
<i>Zichen Kong, Zhongyi Wu, Xiyuan Tang, Yuan Wang</i>	
An Environment-Aware Verification Framework for LLM-Generated Robot Control Programs	2090
<i>Zhanshang Nie, Wenbo Wang, Xuanming Liu, Yue Zhang, Zhendong Chen, Zirui Wang, Kai Huang Shuai Zhao</i>	
Kolmogorov–Arnold Networks for Autonomous Driving: A Hardware-in-the-Loop Comparison with Conventional Deep Neural Networks	2097
<i>Chaoran Yuan, Fadi Kurdahi</i>	
STAR: High-DoF Robotic Manipulation for Memory-Constrained NN Accelerator	2104
<i>Jhao-Ying Chen, Wen Sheng Lim, Tei-Wei Kuo, Yuan-Hao Chang</i>	
A Real-Time Robotic Diffusion Policy Accelerator Exploiting Self- and Cross-Guided Modal Similarity	2111
<i>Boju Chen, Xiaoyu Feng, Junyan Lin, Chen Tang, Huazhong Yang, Yongpan Liu</i>	
TEE-based On-demand Key Distribution for Hierarchical In-Vehicle Zonal Architecture	2118
<i>Wonseok Song, Sanghoon Jeon, Jong-Chan Kim</i>	
FALCON: A Fast and Low-Power Current-Mode Near-Sensor-Computing Architecture for Real-Time Edge Visual Processing	2125
<i>Liang Zhang, Jing Kou, Jinyao Mi, Yang Liu, Junda Zhao, Junzhan Liu, Wang Kang</i>	
LogHD: Robust Compression of Hyperdimensional Classifiers via Logarithmic Class-Axis Reduction	2128
<i>Sanggeon Yun, Hyunwoo Oh, Ryoza Masukawa, Pietro Mercati, Nathaniel D. Bastian, Mohsen Imani</i>	
Boosting the Performance of Tree-Based Speculative Decoding of LLMs on FPGAs	2135
<i>Tielong Liu, Gang Li, Zitao Mo, Zeyu Zhu, Minnan Pei, Jian Cheng</i>	
Equicore: Accelerating Clebsch-Gordan Tensor Product of Equivariant Neural Networks on FPGA	2142
<i>Shidi Tang, Chuanzhao Zhang, Ruiqi Chen, Yuxuan Lv, Bruno Da Silva, Ming Ling</i>	
QUILL: An Algorithm–Architecture Co-Design for Cache-Local Deformable Attention	2149
<i>Hyunwoo Oh, Hanning Chen, Sanggeon Yun, Yang Ni, Wenjun Huang, Tamoghno Das, Suyeon Jang Mohsen Imani</i>	

FLARE: Finetuning ReLU With FIRE for Efficient Long-Context Inference	2156
<i>Michael Moffatt, Junyi Luo, Haoran Cheng, Qilong Wang, Xinting Jiang, Guanchen Tao, Shiwei Liu</i>	
<i>Kauna Lei, Gregory Kielian, Mehdi Saligane</i>	
CD-PIM: A High-Bandwidth and Compute-Efficient LPDDR5-Based PIM for Low-Batch LLM Acceleration on Edge-Device	2163
<i>Ye Lin, Chao Fang, Xiaoyong Song, Qi Wu, Anying Jiang, Yichuan Bai, Li Du</i>	
Identifying Hardware Optimizations for Neural Network Inference using Virtual Prototypes	2170
<i>Jan Zielasko, Rolf Drechsler</i>	
Towards Bit-Shareable Inference on Microcontrollers	2177
<i>Charalampos Bezaits, Yaman Umuroglu, Di Liu, Magnus Sjölander</i>	
Scalable Symbolic Reasoning with Matrix-Based Brain-Inspired Representations and Vector-Space Acceleration	2180
<i>William Youngwoo Chung, Hyunwoo Oh, Hamza Errahmouni Barkam, Calvin Yeung, Mohsen Imani</i>	
Late Breaking Results: Input Loss Curvature as a Predictor of Sample Vulnerability to Hardware-Noise	2183
<i>Deepika Sharma, Deepak Ravikumar, Chih-Hsing Ho, Sangamesh D. Kodge, Kaushik Roy</i>	
Late Breaking Results: Algorithm-Hardware Co-Design of a Sparsity-Aware Dense-Sparse Scheme for DNN Accelerators	2186
<i>Yueting Li, Zhen Li, Terry Tao Ye, Weisheng Zhao</i>	
Late Breaking Results: Conversion of Neural Networks into Logic Flows for Edge Computing	2188
<i>Daniel Stein, Shaoyi Huang, Rolf Drechsler, Bing Li, Grace Li Zhang</i>	
Late Breaking Results: POSEiDON: Pose Estimation in Dynamic On-device Networks via Hyperdimensional Computing	2191
<i>Colin Dupuis, Emilien Meyer, Abu Kaisar Mohammad Masum, Serkan Aygun</i>	
Late Breaking Results: Uncovering the Limits of ECCs in Vision Transformers and a Zero-Cost Reliability Enhancement	2194
<i>Mohammad Hasan Ahmadilivani, Marten Roots, Marco Restifo, Sven-Markus Loorits, Luca Di Mauro</i>	
<i>Jaan Raik</i>	
Late Breaking Results: Quamba-SE: Soft-edge Quantizer for Activations in State Space Models	2197
<i>Yizhi Chen, Ahmed Hemani</i>	
Late Breaking Results: SP-HD: Stochastic Projection-Based HyperDimensional Architecture for Near-Sensor Image Classification	2200
<i>Ahmed Mamdouh, Sabrina Hassan Moon, Abu Kaisar Mohammad Masum, Emilien J. Meyer, Serkan Aygun</i>	
<i>Dayane Reis</i>	
Late Breaking Results: Adaptive Ensembles of Dynamic DNNs for Collaborative Edge Inference	2203
<i>Mingyu Hu, Amit Kumar Singh, Jonathon Hare, Geoff V. Merrett</i>	
Focus Session Paper: The MQT Compiler Collection : A Blueprint for a Future-Proof Quantum-Classical Compilation Framework	2206
<i>Lukas Burgholzer, Daniel Haag, Yannick Stade, Damian Rovara, Patrick Hopf, Robert Wille</i>	
Near-Memory Architecture for Threshold-Ordinal Surface-Based Corner Detection of Event Cameras	2213
<i>Hongyang Shang, An Guo, Shuai Dong, Junyi Yang, Ye Ke, Arindam Basu</i>	

PRISM: A Locality-Aware Near-Memory Processing Framework for Scalable Triangle Counting	2220
<i>Shangdong Zhang, Xueyan Wang, Yier Jin</i>	
Leveraging Recurrent Patterns in Graph Accelerators	2227
<i>Masoud Rahimi, Sébastien Le Beux</i>	
CHIME: Chiplet-based Heterogeneous Near-Memory Acceleration for Edge Multimodal LLM Inference	2234
<i>Yanru Chen, Runyang Tian, Yue Pan, Zheyu Li, Weihong Xu, Tajana Rosing</i>	
DreamRAM: A Fine-Grained Configurable Design Space Modeling Tool for Custom 3D Die-Stacked DRAM	2241
<i>Victor Cai, Jennifer Zhou, David Brooks, Gu-Yeon Wei</i>	
CIM-Tuner: Balancing the Compute and Storage Capacity of SRAM-CIM Accelerator via Hardware-mapping Co-exploration	2248
<i>Jinwu Chen, Yuhui Shi, He Wang, Zhe Jiang, Jun Yang, Xin Si, Zhenhua Zhu</i>	
HINT: A Hybrid SRAM-MRAM Compute-In-Memory with INput-aware Skipping SAR-ADC for Energy Efficient Ternary LLMs	2255
<i>Jaebeom Park, Seung-Eon Hwang, Jongsun Park</i>	
ReBIT: A ReRAM-Based In-Situ Training Accelerator with Robustness Against Stochasticity	2262
<i>Peng Dang, Wei Wang, Yintao He, Huawei Li</i>	
Re-RIS: A Reconfigurable 3D RRAM In-Sensor Architecture for Low-Latency Machine Vision	2264
<i>Shiyang Li, Lixia Han, Siyuan Chen, Lifeng Liu, Peng Huang</i>	
Enhanced CXL Pooled Memory System for Scalable AI via Embedding Access Prediction	2267
<i>Jongho Park, Hoyeon Lee, Seohyun Kim, Minho Ha, Byungil Koh, Jungmin Choi, Yeseong Kim</i>	
FARM: Fast Acceleration of Random forests via in-Memory traversal	2274
<i>Aymen Ahmed, Valeria Bertacco</i>	
Endor: Exploit Nearly-Decode-Only Opportunities of LLM Reasoning on Near-Memory Architecture	2281
<i>Jun Liu, Tianlang Zhao, Shiyi Liu, Jiancai Ye, Lin Li, Zhen Yu, Li Ding, Hao Zhou, Zhenhua Zhu</i>	
<i>Xuefei Ning, Yuan Xie, Yu Wang, Guohao Dai</i>	
Khepri: Crystallizing TAGE for Memory Efficient Prewarm in Serverless Computing	2288
<i>Zengshi Wang, Zhiyuan Zhang, Zhuoyuan Yang, Kanheng Jiang, Chao Fu, Jun Han</i>	
MACAM: A Flexible Computing-in-Memory Accelerator for Sparse Matrix-Dense Vector Multiplication	2295
<i>Xiaoyu Zhang, Rui Liu, Zerun Li, Yinhe Han, Xiaoming Chen</i>	
Zion: A Comprehensive, Adaptive, and Lightweight Hardware Prefetcher	2302
<i>Vadim Biryukov, Xiaoyang Lu, Zirui Liu, Kaixiong Zhou, Xian-He Sun</i>	
Boosting LLC Bandwidth Utilization in GPUs through Adaptive Fine-Grained Data Migration	2309
<i>Jihun Yoon, Sungbin Jang, Seokin Hong</i>	
A High-Performance Neural Rendering Accelerator Based on Novel Multi-Level Ray Scheduling and Dual-Process Backend	2316
<i>Wenkai Zhou, Yuefeng Zhang, Cheng Zhang, Binzhe Yuan, Junsheng Chen, Luntian Zhang, Xiangyu Zhang</i>	
<i>Pingqiang Zhou, Jingyi Yu, Xin Lou</i>	

Data Distribution-Aware Analog/Digital Conversion Strategy for Energy-Efficient Memristive In-Situ Accelerators	2319
<i>Taoming Lei, Heng Zhou, Bing Wu, Wei Tong, Dan Feng</i>	
RATuner: Retrieval-Augmented VLSI Flow Design Parameter Tuning Framework	2322
<i>Peng Xu, Ziyang Yu, Yuan Pu, Xinyun Zhang, Donger Luo, Hao Geng, Siyuan Xu, Tsung-Yi Ho, Bei Yu</i>	
SATA: Sparsity-Aware Scheduling for Selective Token Attention	2329
<i>Zhenkun Fan, Zishen Wan, Che-Kai Liu, Ashwin Sanjay Lele, Win-San Khwa, Bo Zhang, Meng-Fan Chang Arijit Raychowdhury</i>	
TT-Edge: A Hardware–Software Co-Design for Energy-Efficient Tensor-Train Decomposition on Edge AI	2336
<i>Hyunseok Kwak, Kyeongwon Lee, Kyeongpil Min, Chaebin Jung, Woojoo Lee</i>	
Smart Imager with Object Detection Exploiting Edge-Frame-Base Processing and Bounding Box Extraction for μ W Power Purely-Harvested Sensor Nodes	2343
<i>Hayate Okuhara, Udari De Alwis, Liu Yue, Karim Ali Ahmed, Massimo Alioto</i>	
Input Sparsity Aware In-Memory Computing Macro Based on SOT-MRAM Multi-Level Cell for Efficient Deep Neural Network Acceleration	2350
<i>Chao Wang, Qihang Gao, Xianzeng Guo, Zhongzhen Tong, Zhaohao Wang, Weisheng Zhao</i>	
Dynamic Voltage, Body Bias and Frequency Scaling for FD-SOI-Based Low-Power Edge Processors	2357
<i>Shrihari Gokulachandran, Navneet Jain, Andreas Gerstlauer</i>	
GLEAM: A Graph-Learning Enhanced Adaptive Metaheuristic for Power-Aware Scheduling on Heterogeneous Cyber-Physical Systems	2364
<i>Amir Hossein Ansari, Mohsen Ansari, Sepideh Safari, Alireza Ejlali, Jörg Henkel</i>	
SA-ANT: Efficient Low-Bit Group-Wise Quantization for Large Language Models via Sign-Asymmetric Adaptive Numeric Type	2371
<i>Xinkuang Geng, Siting Liu, Hui Wang, Jie Han, Honglan Jiang</i>	
Efficient Down-sampling in Hybrid Neural Networks using Adversarial Autoencoders	2378
<i>Jonghyeon Nam, JoonSeok Kim, Eunji Kwon, Seokhyeong Kang</i>	
Evaluation of Thermal and Power integrity and its Impact on Performance for 3D Memory-on-Logic CPUs with FSPDN and BSPDN	2381
<i>Yumeng Wang, Xincheng Liu, Hu Zhou, Linqiu Wang, Haolan Yang, Zhuojun Chen, Zhiyong Zhang Lianmao Peng, Rongmei Chen</i>	
Focus Session: Autonomous Systems Dependability in the era of AI: Design Challenges in Safety, Security, Reliability and Certification	2384
<i>Behnaz Ranjbar, Kirankumar Raveendiran, Sudeep Pasricha, Samarjit Chakraborty, Cecilia Carbonelli Akash Kumar</i>	
PREFACE: Proactive Re-executions for Fault-aware Mixed-criticality Environments	2391
<i>Hwisoo So, Byeonggil Jun, Chanhee Lee, Hokeun Kim, Aviral Shrivastava</i>	
Automatic Extraction of Timing Models for WCET Estimation From a High-Level Synthesis Flow	2398
<i>Thomas Feuilletin, Dylan Leothaud, Simon Rokicki, Steven Derrien, Isabelle Puaut</i>	
Simultaneous Multithreading and Common-Period Sporadic Tasks in Hard Real-Time	2405
<i>Sims Hill Osborne</i>	

Enabling Ultra-Reliable Memories: A Practical Framework for Zero Mis-correction SEC-DED-DAEC Codes for Safety-Critical Systems	2412
<i>Guixiang Chen, Sheng Liu, Bo Yuan, Yang Guo</i>	
Efficient Warpage Simulation of Complex 2.5-D/3-D IC Structures with Novel Meshing Algorithm and Layerwise Plate Theory	2419
<i>Tianxiang Zhu, Qipan Wang, Yibo Lin, Runsheng Wang</i>	
Exploiting Variable-Dimensional LDPC Coding to Improve NAND Flash Memory System Performance	2426
<i>Meng Zhang, Wei Li, Yangyi Li, Tianwei Gui, Changsheng Xie, Fei Wu</i>	
Topology-Aware Circuit Breaking on Critical Paths in Microservice Systems	2433
<i>Lin Wang, Xin Li, Yanling Bu, Tianhao Zhang, Meiyang Teng, Yanchao Zhao</i>	
SSALDPC: A Syndrome-Sum Based Adaptive LDPC Decoding Scheme for NAND Flash Memory	2440
<i>Lanlan Cui, Fei Wu, Yunlong He, Kun Jiang, Yeqiu Xiao, Renzhi Xiao, Changsheng Xie</i>	
Preemption Threshold Assignment to Improve Schedulability under Memory Constraints	2447
<i>Thilanka Thilakasiri, Matthias Becker</i>	
APEX: Integer-only Non-linear Function Approximation for Efficient Cross-Modal Inference	2450
<i>Peihuan Ni, Zitao Mo, Tielong Liu, Hongli Wen, Zeyu Zhu, Minnan Pei, Junwen Si, Weifan Guan, Peisong Wang, Qinghao Hu, Gang Li, Jian Cheng</i>	
Unary Positional System: Flexible Balance of Hardware Area and Performance	2457
<i>Zeshi Liu, Zheng Weng, Ruijie Tan, Guangming Tang, Haihang You</i>	
DS-CIM: Digital Stochastic Computing-In-Memory Featuring Accurate OR-Accumulation via Sample Region Remapping for Edge AI Models	2464
<i>Kunming Shao, Liang Zhao, Jiangnan Yu, Zhipeng Liao, Xiaomeng Wang, Yi Zou, Tim Kwang-Ting Cheng, Chi-Ying Tsui</i>	
LUT-APP: Dynamic-Precision LUT-based Approximation Unifying Non-Linear Operations in Transformers	2471
<i>Seokkyu Yoon, Namjoon Kim, Hyun Kim</i>	
Learning to Approximate: Circuit Learning and Deep Reinforcement Learning for Approximate Logic Synthesis with an Error Rate Guarantee	2478
<i>Chi-Wei Chen, Yi-Ting Li, Wuqian Tang, Yung-Chih Chen, Jian-Meng Yang, Chun-Yao Wang</i>	
Voltage Aware Approximate CGRA Synthesis for Energy Efficient DNN Inference	2485
<i>Georgios Alexandris, Panagiotis Chaidos, Alexis Maras, Barry de Bruin, Manil Dev Gomony, Henk Corporaal, Dimitrios Soudris, Sotirios Xydis</i>	
Bridging the Power Estimation Gap: A GNN-Based Prediction Model for Approximate Logic Synthesis	2492
<i>Fuxuan Li, Ao Liu, Siting Liu, Hui Wang, Jie Han, Honglan Jiang</i>	
HAWX: A Hardware-Aware FrameWork for Fast and Scalable ApproXimation of DNNs	2499
<i>Samira Nazari, Mohammad Saeed Almasi, Mahdi Taheri, Ali Azarpeyvand, Ali Mokhtari, Ali Mahani, Christian Herglotz</i>	
ARCSyn: Aging-Aware Accuracy-Reconfigurable Logic Synthesis	2506
<i>Ruicheng Dai, Feiyang Shu, Pengpeng Ren, Runsheng Wang, Weikang Qian</i>	

Towards Input-Distribution-Aware Approximate Multiplier Generation for CNNs	2509
<i>Alessandro Buccolini, Marco Biasion, Rodrigo Otoni, George A. Constantinides, Laura Pozzi</i>	
CISim: ISA-Agnostic Custom Instruction Simulation for General-Purpose Processor	2512
<i>Xiaoyu Hao, Sen Zhang, Liang Qiao, Jun Shi, Junshi Chen, Hong An</i>	
Machine Learning-Driven Early Performance Prediction Framework for Accelerated Microarchitecture Simulation	2519
<i>Aiden Stickney, Osvaldo Castro, Aaron Chan, Paul V. Gratz, Jiang Hu, Aakash Tyagi Jered Dominguez-Trujillo, Galen Shipman, Kevin Sheridan</i>	
VeriRepair: Toward Reliable LLM-Based RTL Repair via CoT-Supervised Multi-Objective Fine-Tuning and Hybrid Retrieval	2526
<i>Lei Peng, Aijiao Cui, Yier Jin</i>	
RAMP: RTL-Level Emulation with Thousand-Core-Scale Parallelism	2533
<i>Weigang Feng, Yijia Zhang, Zekun Wang, Zhengyang Wang, Yi Wang, Peijun Ma, Ningyi Xu</i>	
GE-LLM: Graph-Enhanced Large Language Models for Efficient Transistor-Level Circuit Simulation	2540
<i>Chao Wang, Dan Niu, Yichao Dong, Dekang Zhang, Changyin Sun, Zhou Jin</i>	
SimFuzz: Similarity-guided Block-level Mutation for RISC-V Processor Fuzzing	2547
<i>Hao Lyu, Jingzheng Wu, Xiang Ling, Yicheng Zhong, Zhiyuan Li, Tianyue Luo</i>	
Enhance Language Model-based Repair for Memory-related Vulnerabilities via Knowledge-and Semantic-guided Analysis	2554
<i>Hao Shen, Ming Hu, Yanxin Yang, Xiaofei Xie, Mingsong Chen</i>	
Multicore Design Verification Directed by Reinforcement Learning	2561
<i>Luiz M. V. Pereira, Diego M. Meditsch, Guilherme O. Campos, Luiz C. V. Dos Santos</i>	
Hetero-ChipletSim: Bridging Chiplet, Interconnect and Packaging Heterogeneity in Multi-Chiplet System Simulation	2568
<i>Xuguang Yuan, Jiangyuan Gu, Qidie Wu, Yang Hu, Shaojun Wei, Shouyi Yin</i>	
CoverAssert: Iterative LLM Assertion Generation Driven by Functional Coverage via Syntax-Semantic Representations	2571
<i>Yonghao Wang, Yang Yin, Hongqin Lyu, Jiaxin Zhou, Zhiteng Chao, Mingyu Shi, Wenchao Ding, Yunlin Du Jing Ye, Tiancheng Wang, Huawei Li</i>	
SuperPhys-Net: A Physics-Informed Super-Resolution Electromagnetic Simulator for Nanophotonic Devices	2574
<i>Yiyang Su, Hao Chen, Guohao Dai, Yuzhe Ma, Yeyu Tong</i>	
Critical-Path-Centric 3D IC Placement for Timing Optimization	2581
<i>Sojung Park, Heechun Park</i>	
SFQ-Based CJoin Gate Implementation for Ultra-Low-Power Brownian Logic Circuits	2588
<i>Soshi Takagi, Masamitsu Tanaka, Koji Inoue, Satoshi Kawakami</i>	
Architecture, Design and Technology Co-optimization for 3D ICs with Advanced BSPDN Considering Power & Thermal Integrity Impact	2595
<i>Hu Zhou, Haolan Yang, Xingcheng Liu, Linqiu Wang, Feifan Xie, Yumeng Wang, Zhuojun Chen Zhiyong Zhang, Lianmao Peng, Rongmei Chen</i>	

3D IC Thermal Management with BEOL Wafer-Scale Sputtered Vertical h-BN	2602
<i>C. Smith, B. Reese, A. Raut, Y.-T. Yang, J.-G. Zhu, T. Srimani</i>	
AgenticTCAD: A LLM-based Multi-Agent Framework for Automated TCAD Code Generation and Device Optimization	2609
<i>Guangxi Fan, Tianliang Ma, Xuguang Sun, Xun Wang, Kain Lu Low, Leilai Shao</i>	
IncreMacro-3D: Incremental Macro Placement for Face-to-Face Stacked Memory-on-Logic 3D ICs	2616
<i>Lancheng Zou, Sing Sen Ye, Shuo Yin, Yuan Pu, Jiayi Jiang, Siting Liu, Yuxuan Zhao, Bei Yu</i>	
SCARLET: A Scalable OPCM-Based Accelerator for Transformer Inference with Tiled Crossbars	2623
<i>Sina Karimi, Guowei Yang, Carlos A. Ríos Ocampo, Ajay Joshi, Ayse K. Coskun</i>	
Braid-ZNS: Leveraging Zone Random Write Area for Efficient In-Storage Compression on ZNS SSDs	2630
<i>Minkyu Choi, Joonseong Hwang, Minjin Park, Seokin Hong</i>	
Focus Session: LLM4PQC – Accurate and Efficient Synthesis of PQC Cores by Feedback-Driven LLMs	2633
<i>Buddhi Perera, Zeng Wang, Weihua Xiao, Mohammed Nabeel, Ozgur Sinanoglu, Johann Knechtel Ramesh Karri</i>	
Optimize edge AI processing through innovative compilation techniques	2640
<i>Shreya Alladi, Alexandre Lopoukhine, Georgios Alexandris, Andrea Nardi-Dei Ravikiran Ravindranath Reddy, Christos P. Lamprakos, Panagiotis Chaidos, Alexis Maras, Alberto Ros Tobias Grosser, Sotirios Xydis, Dimitrios Soudris, Marc Geilen, Sander Stuijk, Henk Corporaal Alexandra Jimborean</i>	
Multi-Partner Project: Quantum-Secure IoT-based Digital Manufacturing Pilot in QUBIP project	2647
<i>Eros Camacho-Ruiz, Pablo Navarro-Torrero, Piedad Brox, Maria Chiara Molteni, Alberto Battistello Davide Bellizia, Agostino Sette, Enrico Bisio, Nicola Tuveri, Enrico Bravi, Francesco Vaccaro Grazia D’Onghia, Andrea Vesco</i>	
Multi-Partner Project: dAIEDGE - A Network of Excellence for Distributed, Trustworthy, Efficient and Scalable AI at the Edge	2653
<i>Alain Pagani, José Cano, Haralampos-G. Stratigopoulos, Aysajan Abidin, Mhd Rashed Al Koutayni Luca Benini, Angelos Bilas, Alessandro Capotondi, Roberto Cavicchioli, Brian Clerkin, Oscar Deniz Margaux Divernois, Baptiste Dupertuis, Dorvan Favre, Giulio Gambardella, Ander Garcia Gangoiti Carlo Augusto Grazia, Dominik Günzel, Jude Haris, Klodjan K. Hidri, Maïck Huguenin, Manal Jammal Paul Kling, Christos Kozanitis, Xavier Lessage, Srikanth Mandapati, Philippe Massonet, Alfio Di Mauro Varesh Mishra, Juan Odriozola, Javier Parra Domínguez, Nuria Pazos, Viviane Potocnik, Miguel de Prado Rohit Prasad, Spyridon Raptis, Gregoire Rebstein, Ignacio Sañudo Olmedo, Mohamed Selim Chinmay Satish Shrivastav, Noelia Vallez, Giorgos Vasiliadis, Micaela Verrucchi, Enrico Vincenzi Damian Vizár, Devendra Vyas, Stefan Wiehle</i>	
Multi-Partner Project: A Holistic and Open-Source Approach to Efficient, Secure and Reliable AI Hardware Deployment in DI-EDAI	2660
<i>Georgios Sotiropoulos, Felix Frombach, Julian Hoefer, Tanja Harbaum, Juergen Becker Henrik Iver Thorøe, Vincent Meyers, Mehdi Tahoori, Zeynep Demirdag, Mohammed Bakr Sikal Hassan Nassar, Heba Khdr, Jörg Henkel, Christopher Wolters, Philipp van Kempen, Johannes Geier Ulf Schlichtmann, Batuhan Sesli, Muhammad Sabih, Jakob Wittmann, Frank Hannig, Jürgen Teich Lukas Steiner, Norbert Wehn, Mohamed Shelkamy Ali, Philipp Schmitz, Wolfgang Kunz, Stefan Koegler Georg Sigl</i>	

Multi-Partner Project: Enhancing Resilience, Efficiency, and Trustworthiness of Edge AI in Safety-Critical Systems (GuardAI)	2667
<i>Antonis Savva, Mehmet Demirel, Yeshwanth Kumar Adimoolam, Rafaella Elia, Alexandros Gkillas Christos Anagnostopoulos, Erion-Vasilis Pikoulis, Amalia Damianou, Charmaine Barker, Daniel Bethell Ahmed Salah Tawfik Ibrahim, Filippo Cugini, Francesco Paolucci, Kyriakos Vlachos, Simos Gerasimou Antonios Lalas, Konstantinos Votis, Aris Lalos, Christos Kyrkou, Theocharis Theocharides</i>	
Multi-Partner Project: Scalable, Ferroelectric-based Accelerators for Energy Efficient Edge AI (Ferro4EdgeAI)	2674
<i>Theofilos Spyrou, Yashvardhan Biyani, Konstantinos Stavrakakis, Rajendra Bishnoi, Said Hamdioui Joel Minguet Lopez, Louise Dumas, Jean Coignus, Denys Ly, Hugo Chazot-Ranquet, Laurent Grenouillet Fabien Grimaud, Simon Martin, Olivier Billoint, Francois Andrieu, Ruben Alcala, Stefan Slesazeck Athira Sunil, Chong Peng, Antoine Cauquil, Rosario Pronsato, Damien Deleruyelle, Cédric Marchand Alberto Bosio, Ian O'Connor, Giulio Urlini, Simon Jeannot, Mohammad Sajedi Alvar Nima Akbari Moghaddam, Thilo Werner, Tony Schenk, Bojun Cheng, Mina Khoei, Lucía Pérez Ramírez EunJin Koh, Somnath Kale, Nicholas Barrett</i>	
Multi-Partner Project: Outcomes of the ICSC Flagship 2 Project on Architectures and Design Methodologies to Accelerate AI Workloads	2681
<i>Cristina Silvano, Fabrizio Ferrandi, Serena Curzel, Daniele Ielmini, Cristian Zambelli Sebastiano Fabio Schifano, Francesco Conti, Angelo Garofalo, Luca Benini, Maurizio Palesi Giuseppe Ascia, Enrico Russo, Fanny Spagnolo, Pasquale Corsonello, Stefania Perri, Fabio Frustaci</i>	
Multi-Partner Project: Efficient Deep Learning Platforms for Next-Generation Embedded Edge-AI Systems	2688
<i>Rajendra Bishnoi, Mohammad Amin Yaldagard, Konstantinos Stavrakakis, Said Hamdioui Kanishkan Vadivel, Pankaj Upadhyay, Nicolas Daniel Rodriguez, Teresa van Dam Sander Steeghs-Turchina, Agathe Archet, Prathamesh Satish Deshpande, Giovanni Grandi, Hana Krichene William Fabre, Fabian Chersi</i>	