

2022 59th ACM/IEEE Design Automation Conference (DAC 2022)

**San Francisco, California, USA
6-10 June 2022**

Pages 1-696



**IEEE Catalog Number: CFP22DAC-POD
ISBN: 979-8-3315-6500-8**

**Copyright © 2022, ACM
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:	CFP22DAC-POD
ISBN (Print-On-Demand):	979-8-3315-6500-8
ISBN (Online):	978-1-4503-9142-9
ISSN:	0738-100X

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

10.1	QuantumNAT: Quantum Noise-Aware Training with Noise Injection, Quantization and Normalization	1
	<i>Hanrui Wang, Jiaqi Gu, Yongshan Ding, Zirui Li, Frederic T. Chong, David Z. Pan, Song Han</i>	
10.2	Optimizing Quantum Circuit Synthesis for Permutations Using Recursion	7
	<i>Cynthia Chen, Bruno Schmitt, Helena Zhang, Lev S. Bishop, Ali Javadi-Abhari</i>	
10.3	A Fast and Scalable Qubit-Mapping Method for Noisy Intermediate-Scale Quantum Computers	13
	<i>Sunghye Park, Daeyeon Kim, Minhyuk Kweon, Jae-Yoon Sim, Seokhyeong Kang</i>	
10.4	Optimizing Quantum Circuit Placement via Machine Learning	19
	<i>Hongxiang Fan, Ce Guo, Wayne Luk</i>	
11.1	HERO: Hessian-Enhanced Robust Optimization for Unifying and Improving Generalization and Quantization Performance	25
	<i>Huanrui Yang, Xiaoxuan Yang, Neil Zhenqiang Gong, Yiran Chen</i>	
11.2	Neural Computation for Robust and Holographic Face Detection	31
	<i>Mohsen Imani, Ali Zakeri, Hanning Chen, TaeHyun Kim, Prathyush Poduval, Hyunsei Lee, Yeseong Kim, Elaheh Sadredini, Farhad Imani</i>	
11.3	FHDnn: Communication Efficient and Robust Federated Learning for AIoT Networks	37
	<i>Rishikanth Chandrasekaran, Kazim Ergun, Jihyun Lee, Dhanush Nanjunda, Jaeyoung Kang, Tajana Rosing</i>	
11.4	ODHD: One-Class Brain-Inspired Hyperdimensional Computing for Outlier Detection	43
	<i>Ruixuan Wang, Xun Jiao, X. Sharon Hu</i>	
12.1	High-Level Synthesis Performance Prediction Using GNNs: Benchmarking, Modeling, and Advancing	49
	<i>Nan Wu, Hang Yang, Yuan Xie, Pan Li, Cong Hao</i>	
12.2	Automated Accelerator Optimization Aided by Graph Neural Networks	55
	<i>Atefeh Sohrabizadeh, Yunsheng Bai, Yizhou Sun, Jason Cong</i>	
12.3	Functionality Matters in Netlist Representation Learning	61
	<i>Ziyi Wang, Chen Bai, Zhuolun He, Guangliang Zhang, Qiang Xu, Tsung-Yi Ho, Bei Yu, Yu Huang</i>	
12.4	EMS: Efficient Memory Subsystem Synthesis for Spatial Accelerators	67
	<i>Liancheng Jia, Yuyue Wang, Jingwen Leng, Yun Liang</i>	
13.1	DA PUF: Dual-State Analog PUF	73
	<i>Jiliang Zhang, Lin Ding, Zhuojun Chen, Wenshang Li, Gang Qu</i>	
13.2	PathFinder: Side Channel Protection through Automatic Leaky Paths Identification and Obfuscation	79
	<i>Haocheng Ma, Qizhi Zhang, Ya Gao, Jiaji He, Yiqiang Zhao, Yier Jin</i>	

13.3	LOCK & ROLL: Deep-Learning Power Side-Channel Attack Mitigation Using Emerging Reconfigurable Devices and Logic Locking	85
	<i>Gaurav Kolhe, Tyler Sheaves, Kevin Immanuel Gubbi, Soheil Salehi, Setareh Rafatirad, Sai Manoj Pudukotai Dinakarrao, Avesta Sasan, Houman Homayoun</i>	
13.4	Efficient Access Scheme for Multi-bank Based NTT Architecture through Conflict Graph	91
	<i>Xiangren Chen, Bohan Yang, Yong Lu, Shouyi Yin, Shaojun Wei, Leibo Liu</i>	
14.1	InfoX: An Energy-Efficient ReRAM Accelerator Design with Information-Lossless Low-Bit ADCs	97
	<i>Yintao He, Songyun Qu, Ying Wang, Bing Li, Huawei Li, Xiaowei Li</i>	
14.2	PHANES: ReRAM-based Photonic Accelerators for Deep Neural Networks	103
	<i>Yinyi Liu, Jiaqi Liu, Yuxiang Fu, Shixi Chen, Jiaxu Zhang, Jiang Xu</i>	
14.3	CP-SRAM: Charge-Pulsation SRAM Marco for Ultra-High Energy-Efficient Computing-in-Memory	109
	<i>He Zhang, Linjun Jiang, Jianxin Wu, Tingran Chen, Junzhan Liu, Wang Kang, Weisheng Zhao</i>	
14.4	CREAM: Computing in ReRAM-assisted Energy and Area-efficient SRAM for Neural Network Acceleration	115
	<i>Liukai Xu, Songyuan Liu, Zhi Li, Dengfeng Wang, Yiming Chen, Yanan Sun, Xueqing Li, Weifeng He, Shi Xu</i>	
15.1	Chiplet Actuary: A Quantitative Cost Model and Multi-Chiplet Architecture Exploration	121
	<i>Yinxiao Feng, Kaisheng Ma</i>	
15.2	PANORAMA: Divide-and-Conquer Approach for Mapping Complex Loop Kernels on CGRA	127
	<i>Dhananjaya Wijerathne, Zhaoying Li, Thilini Kaushalya Bandara, Tulika Mitra</i>	
15.3	A Fast Parameter Tuning Framework via Transfer Learning and Multi-objective Bayesian Optimization	133
	<i>Zheng Zhang, Tinghuan Chen, Jiaxin Huang, Meng Zhang</i>	
15.4	PriMax: Maximizing DSL Application Performance with Selective Primitive Acceleration	139
	<i>Nicholas Wendt, Todd Austin, Valeria Bertacco</i>	
16.1	Accelerating and Pruning CNNs for Semantic Segmentation on FPGA	145
	<i>Pierpaolo Mori, Manoj-Rohit Vemparala, Nael Fafous, Saptarshi Mitra, Sreetama Sarkar, Alexander Frickenstein, Lukas Frickenstein, Domenik Helms, Naveen Shankar Nagaraja, Walter Stechele, Claudio Passerone</i>	
16.2	SoftSNN: Low-Cost Fault Tolerance for Spiking Neural Network Accelerators under Soft Errors	151
	<i>Rachmad Vidya Wicaksana Putra, Muhammad Abdullah Hanif, Muhammad Shafique</i>	
16.3	A Joint Management Middleware to Improve Training Performance of Deep Recommendation Systems with SSDs	157
	<i>Chun-Feng Wu, Carole-Jean Wu, Gu-Yeon Wei, David Brooks</i>	
16.4	The Larger the Fairer? Small Neural Networks Can Achieve Fairness for Edge Devices	163
	<i>Yi Sheng, Junhuan Yang, Yawen Wu, Kevin Mao, Yiyu Shi, Jingtong Hu, Weiwen Jiang, Lei Yang</i>	
17.1	SCAIE-V: An Open-Source SCAlable Interface for ISA Extensions for RISC-V Processors	169
	<i>Mihaela Damian, Julian Oppermann, Christoph Spang, Andreas Koch</i>	
17.2	A Scalable Symbolic Simulation Tool for Low Power Embedded Systems	175
	<i>Subhash Sethumurugan, Shashank Hegde, Hari Cherupalli, John Sartori</i>	

17.3	Designing Critical Systems with Iterative Automated Safety Analysis	181
	<i>Ran Wei, Zhe Jiang, Xiaoran Guo, Haitao Mei, Athanasios Zolotas, Tim Kelly</i>	
18.1	Efficient Ensembles of Graph Neural Networks	187
	<i>Amrit Nagarajan, Jacob R. Stevens, Anand Raghunathan</i>	
18.2	Sign Bit is Enough: A Learning Synchronization Framework for Multi-hop All-reduce with Ultimate Compression	193
	<i>Feijie Wu, Shiqi He, Song Guo, Zhihao Qu, Haozhao Wang, Weihua Zhuang, Jie Zhang</i>	
18.3	GLite: A Fast and Efficient Automatic Graph-Level Optimizer for Large-Scale DNNs	199
	<i>Jiaqi Li, Min Peng, Qingan Li, Meizheng Peng, Mengting Yuan</i>	
18.4	Contrastive Quant: Quantization Makes Stronger Contrastive Learning	205
	<i>Yonggan Fu, Qixuan Yu, Meng Li, Xu Ouyang, Vikas Chandra, Yingyan Lin</i>	
19.1	Serpens: A High Bandwidth Memory Based Accelerator for General-Purpose Sparse Matrix-Vector Multiplication	211
	<i>Linghao Song, Yuze Chi, Licheng Guo, Jason Cong</i>	
19.2	An Energy-Efficient Seizure Detection Processor Using Event-Driven Multi-Stage CNN Classification and Segmented Data Processing with Adaptive Channel Selection	217
	<i>Jiahao Liu, Zirui Zhong, Yong Zhou, Hui Qiu, Jianbiao Xiao, Jiajing Fan, Zhaomin Zhang, Sixu Li, Yiming Xu, Siqi Yang, Weiwei Shan, Shuisheng Lin, Liang Chang, Jun Zhou</i>	
19.3	PatterNet: Explore and Exploit Filter Patterns for Efficient Deep Neural Networks	223
	<i>Behnam Khaleghi, Uday Mallappa, Duygu Yaldiz, Haichao Yang, Monil Shah, Jaeyoung Kang, Tajana Rosing</i>	
19.4	E2SR: An End-to-End Video CODEC Assisted System for Super Resolution Acceleration	229
	<i>Zhuoran Song, Zhongkai Yu, Naifeng Jing, Xiaoyao Liang</i>	
20.1	MATCHA: A Fast and Energy-Efficient Accelerator for Fully Homomorphic Encryption over the Torus	235
	<i>Lei Jiang, Qian Lou, Nrushad Joshi</i>	
20.2	VirTEE: A Full Backward-Compatible TEE with Native Live Migration and Secure I/O	241
	<i>Jianqiang Wang, Pouya Mahmoody, Ferdinand Brasser, Patrick Jauernig, Ahmad-Reza Sadeghi, Donghui Yu, Dahan Pan, Yuanyuan Zhang</i>	
20.3	Apple vs. EMA: Electromagnetic Side Channel Attacks on Apple CoreCrypto	247
	<i>Gregor Haas, Aydin Aysu</i>	
21.1	Algorithm/Architecture Co-Design for Energy-Efficient Acceleration of Multi-Task DNN	253
	<i>Jaekang Shin, Seungkyu Choi, Jongwoo Ra, Lee-Sup Kim</i>	
21.2	EBSP: Evolving Bit Sparsity Patterns for Hardware-Friendly Inference of Quantized Deep Neural Networks	259
	<i>Fangxin Liu, Wenbo Zhao, Zongwu Wang, Yongbiao Chen, Zhezhi He, Naifeng Jing, Xiaoyao Liang, Li Jiang</i>	
21.3	A Time-to-first-spike Coding and Conversion Aware Training for Energy-Efficient Deep Spiking Neural Network Processor Design	265
	<i>Dongwoo Lew, Kyungchul Lee, Jongsun Park</i>	

21.4	XMA: A Crossbar-aware Multi-task Adaption Framework via Shift-based Mask Learning Method	271
	<i>Fan Zhang, Li Yang, Jian Meng, Jae-sun Seo, Yu Cao, Deliang Fan</i>	
21.5	SWIM: Selective Write-Verify for Computing-in-Memory Neural Accelerators	277
	<i>Zheyu Yan, Xiaobo Sharon Hu, Yiyu Shi</i>	
22.1	Enabling Efficient Deep Convolutional Neural Network-based Sensor Fusion for Autonomous Driving	283
	<i>Xiaoming Zeng, Zhendong Wang, Yang Hu</i>	
22.2	Zhuyi: Perception Processing Rate Estimation for Safety of Autonomous Vehicles	289
	<i>Yu-Shun Hsiao, Siva Kumar Sastry Hari, Michal Filipiuk, Timothy Tsai, Michael B. Sullivan, Vijay Janapa Reddi, Vasu Singh, Stephen W. Keckler</i>	
22.3	Processing-in-SRAM Acceleration for Ultra-Low Power Visual 3D Perception	295
	<i>Yuquan He, Songyun Qu, Gangliang Lin, Cheng Liu, Lei Zhang, Ying Wang</i>	
22.4	Response Time Analysis for Dynamic Priority Scheduling in ROS2	301
	<i>Abdullah Al Arafat, Sudharsan Vaidhun, Kurt M. Wilson, Jinghao Sun, Zhishan Guo</i>	
22.5	Voltage Prediction of Drone Battery Reflecting Internal Temperature	307
	<i>Jiwon Kim, Seunghyeok Jeon, Jaehyun Kim, Hojung Cha</i>	
23.1	A Near-Storage Framework for Boosted Data Preprocessing of Mass Spectrum Clustering	313
	<i>Weihong Xu, Jaeyoung Kang, Tajana Rosing</i>	
23.2	MetaZip: A High-throughput and Efficient Accelerator for DEFLATE	319
	<i>Ruihao Gao, Xueqi Li, Yewen Li, Xun Wang, Guangming Tan</i>	
23.3	Enabling Fast Uncertainty Estimation: Accelerating Bayesian Transformers via Algorithmic and Hardware Optimizations	325
	<i>Hongxiang Fan, Martin Ferianc, Wayne Luk</i>	
23.4	Eventor: An Efficient Event-Based Monocular Multi-View Stereo Accelerator on FPGA Platform	331
	<i>Mingjun Li, Jianlei Yang, Yingjie Qi, Meng Dong, Yuhao Yang, Runze Liu, Weitao Pan, Bei Yu, Weisheng Zhao</i>	
23.5	GEML: GNN-Based Efficient Mapping Method for Large Loop Applications on CGRA	337
	<i>Mingyang Kou, Jun Zeng, Boxiao Han, Fei Xu, Jiangyuan Gu, Hailong Yao</i>	
23.6	Mixed-granularity Parallel Coarse-grained Reconfigurable Architecture	343
	<i>Jinyi Deng, Linyun Zhang, Lei Wang, Jiawei Liu, Kexiang Deng, Shibin Tang, Jiangyuan Gu, Boxiao Han, Fei Xu, Leibo Liu, Shaojun Wei, Shouyi Yin</i>	
24.1	GuardNN: Secure Accelerator Architecture for Privacy-Preserving Deep Learning	349
	<i>Weizhe Hua, Muhammad Umar, Zhiru Zhang, G. Edward Suh</i>	
24.2	SRA: A Secure ReRAM-Based DNN Accelerator	355
	<i>Lei Zhao, Youtao Zhang, Jun Yang</i>	
24.3	ABNN2: Secure Two-party Arbitrary-Bitwidth Quantized Neural Network Predictions	361
	<i>Liyan Shen, Ye Dong, Binxing Fang, Jinqiao Shi, Xuebin Wang, Shengli Pan, Ruisheng Shi</i>	

24.4	Adaptive Neural Recovery for Highly Robust Brain-like Representation	367
	<i>Prathyush Poduval, Yang Ni, Yeseong Kim, Kai Ni, Raghavan Kumar, Rossario Cammarota, Mohsen Imani</i>	
24.5	Efficiency Attacks on Spiking Neural Networks	373
	<i>Sarada Krithivasan, Sanchari Sen, Nitin Rathi, Kaushik Roy, Anand Raghunathan</i>	
25.1	L-QoCo: Learning to Optimize Cache Capacity Overloading in Storage Systems	379
	<i>Ji Zhang, Xijun Li, Xiyao Zhou, Mingxuan Yuan, Zhuo Cheng, Keji Huang, Yifan Li</i>	
25.2	Pipette: Efficient Fine-Grained Reads for SSDs	385
	<i>Shuhan Bai, Hu Wan, Yun Huang, Xuan Sun, Fei Wu, Changsheng Xie, Hung-Chih Hsieh, Tei-Wei Kuo, Chun Jason Xue</i>	
25.3	CDB: Critical Data Backup Design for Consumer Devices with High-Density Flash based Hybrid Storage	391
	<i>Longfei Luo, Dingcui Yu, Liang Shi, Chuanming Ding, Changlong Li, Edwin H.-M. Sha</i>	
25.4	SS-LRU: A Smart Segmented LRU Caching	397
	<i>Chunhua Li, Man Wu, Yuhan Liu, Ke Zhou, Ji Zhang, Yunqing Sun</i>	
25.5	NobLSM: An LSM-tree with Non-blocking Writes for SSDs	403
	<i>Haoran Dang, Chongnan Ye, Yanpeng Hu, Chundong Wang</i>	
25.6	TailCut: Improving Performance and Lifetime of SSDs Using Pattern-Aware State Encoding	409
	<i>Jaeyong Lee, Myungsuk Kim, Wonil Choi, Sanggu Lee, Jihong Kim</i>	
26.1	HIMap: A Heuristic and Iterative Logic Synthesis Approach	415
	<i>Xing Li, Lei Chen, Fan Yang, Mingxuan Yuan, Hongli Yan, Yupeng Wan</i>	
26.2	Improving LUT-Based Optimization for ASICs	421
	<i>Walter Lau Neto, Luca Amarú, Vinicius Possani, Patrick Vuillod, Jiong Luo, Alan Mishchenko, Pierre-Emmanuel Gaillardon</i>	
26.3	NovelRewrite: Node-Level Parallel AIG Rewriting	427
	<i>Shiju Lin, Jinwei Liu, Tianji Liu, Martin D.F. Wong, Evangeline F.Y. Young</i>	
26.4	Search Space Characterization for Approximate Logic Synthesis	433
	<i>Linus Witschen, Tobias Wiersema, Lucas Reuter, Marco Platzner</i>	
26.5	SEALS: Sensitivity-driven Efficient Approximate Logic Synthesis	439
	<i>Chang Meng, Xuan Wang, Jiajun Sun, Sijun Tao, Wei Wu, Zhihang Wu, Leibin Ni, Xiaolong Shen, Junfeng Zhao, Weikang Qian</i>	
26.6	Beyond Local Optimality of Buffer and Splitter Insertion for AQFP Circuits	445
	<i>Siang-Yun Lee, Heinz Riener, Giovanni De Micheli</i>	
27.1	NAX: Neural Architecture and Memristive Xbar based Accelerator Co-design	451
	<i>Shubham Negi, Indranil Chakraborty, Aayush Ankit, Kaushik Roy</i>	
27.2	MC-CIM: A Reconfigurable Computation-in-Memory for Efficient Stereo Matching Cost Computation	457
	<i>Zhiheng Yue, Yabing Wang, Leibo Liu, Shaojun Wei, Shouyi Yin</i>	

27.3	iMARS: An In-Memory-Computing Architecture for Recommendation Systems	463
	<i>Mengyuan Li, Ann Franchesca Laguna, Dayane Reis, Xunzhao Yin, Michael Niemier, X. Sharon Hu</i>	
27.4	ReGNN: A ReRAM-based Heterogeneous Architecture for General Graph Neural Networks	469
	<i>Cong Liu, Haikun Liu, Hai Jin, Xiaofei Liao, Yu Zhang, Zhuohui Duan, Jiahong Xu, Huize Li</i>	
28.1	You Only Search Once: On Lightweight Differentiable Architecture Search for Resource-Constrained Embedded Platforms	475
	<i>Xiangzhong Luo, Di Liu, Hao Kong, Shuo Huai, Hui Chen, Weichen Liu</i>	
28.2	EcoFusion: Energy-Aware Adaptive Sensor Fusion for Efficient Autonomous Vehicle Perception	481
	<i>Arnav Vaibhav Malawade, Trier Mortlock, Mohammad Abdullah Al Faruque</i>	
28.3	Human Emotion Based Real-time Memory and Computation Management on Resource-Limited Edge Devices	487
	<i>Yijie Wei, Zhiwei Zhong, Jie Gu</i>	
29.1	Hierarchical Memory-Constrained Operator Scheduling of Neural Architecture Search Networks	493
	<i>Zihan Wang, Chengcheng Wan, Yuting Chen, Ziyi Lin, He Jiang, Lei Qiao</i>	
29.2	MIME: Adapting a Single Neural Network for Multi-task Inference with Memory-efficient Dynamic Pruning	499
	<i>Abhiroop Bhattacharjee, Yeshwanth Venkatesha, Abhishek Moitra, Priyadarshini Panda</i>	
29.3	Sniper: Cloud-Edge Collaborative Inference Scheduling with Neural Network Similarity Modeling	505
	<i>Weihong Liu, Jiawei Geng, Zongwei Zhu, Jing Cao, Zirui Lian</i>	
29.4	LPCA: Learned MRC Profiling based Cache Allocation for File Storage Systems	511
	<i>Yibin Gu, Yifan Li, Hua Wang, Li Liu, Ke Zhou, Wei Fang, Gang Hu, Jinhu Liu, Zhuo Cheng</i>	
30.1	Equivalence Checking Paradigms in Quantum Circuit Design: A Case Study	517
	<i>Tom Peham, Lukas Burgholzer, Robert Wille</i>	
30.2	Accurate BDD-based Unitary Operator Manipulation for Scalable and Robust Quantum Circuit Verification	523
	<i>Chun-Yu Wei, Yuan-Hung Tsai, Chiao-Shan Jhang, Jie-Hong R. Jiang</i>	
30.3	Handling Non-Unitaries in Quantum Circuit Equivalence Checking	529
	<i>Lukas Burgholzer, Robert Wille</i>	
30.4	A Bridge-based Algorithm for Simultaneous Primal and Dual Defects Compression on Topologically Quantum-error-corrected Circuits	535
	<i>Wei-Hsiang Tseng, Yao-Wen Chang</i>	
31.1	FaSe: Fast Selective Flushing to Mitigate Contention-based Cache Timing Attacks	541
	<i>Tuo Li, Sri Parameswaran</i>	
31.2	Conditional Address Propagation: An Efficient Defense Mechanism against Transient Execution Attacks	547
	<i>Peinan Li, Rui Hou, Lutan Zhao, Yifan Zhu, Dan Meng</i>	
31.3	Timed Speculative Attacks exploiting Store-to-Load Forwarding bypassing Cache-based Countermeasures	553
	<i>Anirban Chakraborty, Nikhilesh Singh, Sarani Bhattacharya, Chester Rebeiro, Debdeep Mukhopadhyay</i>	

31.4	DARPT: Defense Against Remote Physical Attack based on TDC in Multi-tenant Scenario	559
	<i>Fan Zhang, Zhiyong Wang, Haoting Shen, Bolin Yang, Qianmei Wu, Kui Ren</i>	
32.1	GNNIE: GNN Inference Engine with Load-balancing and Graph-specific Caching	565
	<i>Sudipta Mondal, Susmita Dey Manasi, Kishor Kunal, Ramprasath S, Sachin S. Sapatnekar</i>	
32.2	SALO: An Efficient Spatial Accelerator Enabling Hybrid Sparse Attention Mechanisms for Long Sequences	571
	<i>Guan Shen, Jieru Zhao, Quan Chen, Jingwen Leng, Chao Li, Minyi Guo</i>	
32.3	NN-LUT: Neural Approximation of Non-Linear Operations for Efficient Transformer Inference	577
	<i>Joonsang Yu, Junki Park, Seongmin Park, Minsoo Kim, Sihwa Lee, Dong Hyun Lee, Jungwook Choi</i>	
32.4	Self Adaptive Reconfigurable Arrays (SARA): Learning Flexible GEMM Accelerator Configuration and Mapping-space Using ML	583
	<i>Ananda Samajdar, Eric Qin, Michael Pellauer, Tushar Krishna</i>	
33.1	Enabling Hard Constraints in Differentiable Neural Network and Accelerator Co-Exploration	589
	<i>Deokki Hong, Kanghyun Choi, Hye Yoon Lee, Joonsang Yu, Noseong Park, Youngsok Kim, Jinho Lee</i>	
33.2	Heuristic Adaptability to Input Dynamics for SpMM on GPUs	595
	<i>Guohao Dai, Guyue Huang, Shang Yang, Zhongming Yu, Hengrui Zhang, Yufei Ding, Yuan Xie, Huazhong Yang, Yu Wang</i>	
33.3	H2H: Heterogeneous Model to Heterogeneous System Mapping with Computation and Communication Awareness	601
	<i>Xinyi Zhang, Cong Hao, Peipei Zhou, Alex Jones, Jingtong Hu</i>	
33.4	PARIS and ELSA: An Elastic Scheduling Algorithm for Reconfigurable Multi-GPU Inference Servers	607
	<i>Yunseong Kim, Yujeong Choi, Minsoo Rhu</i>	
34.1	Pursuing More Effective Graph Spectral Sparsifiers via Approximate Trace Reduction	613
	<i>Zhiqiang Liu, Wenjian Yu</i>	
34.2	Accelerating Nonlinear DC Circuit Simulation with Reinforcement Learning	619
	<i>Zhou Jin, Haojie Pei, Yichao Dong, Xiang Jin, Xiao Wu, Wei W. Xing, Dan Niu</i>	
34.3	An Efficient Yield Optimization Method for Analog Circuits via Gaussian Process Classification and Varying-Sigma Sampling	625
	<i>Xiaodong Wang, Changhao Yan, Fan Yang, Dian Zhou, Xuan Zeng</i>	
35.1	Partition and Place Finite Element Model on Wafer-Scale Engine	631
	<i>Jinwei Liu, Xiaopeng Zhang, Shiju Lin, Xinshi Zang, Jingsong Chen, Bentian Jiang, Martin D.F. Wong, Evangeline F.Y. Young</i>	
35.2	CNN-inspired Analytical Global Placement for Large-scale Heterogeneous FPGAs	637
	<i>Huimin Wang, Xingyu Tong, Chenyue Ma, Runming Shi, Jianli Chen, Kun Wang, Jun Yu, Yao-Wen Chang</i>	
35.3	High-performance Placement for Large-scale Heterogeneous FPGAs with Clock Constraints	643
	<i>Ziran Zhu, Yangjie Mei, Zijun Li, Jingwen Lin, Jianli Chen, Jun Yang, Yao-Wen Chang</i>	
35.4	Multi-Electrostatic FPGA Placement Considering SLICEL-SLICEM Heterogeneity and Clock Feasibility	649
	<i>Jing Mai, Yibai Meng, Zhixiong Di, Yibo Lin</i>	

36.1	QOC: Quantum On-Chip Training with Parameter Shift and Gradient Pruning	655
	<i>Hanrui Wang, Zirui Li, Jiaqi Gu, Yongshan Ding, David Z. Pan, Song Han</i>	
36.2	Memory-Efficient Training of Binarized Neural Networks on the Edge	661
	<i>Mikail Yayla, Jian-Jia Chen</i>	
36.3	DeepGate: Learning Neural Representations of Logic Gates	667
	<i>Min Li, Sadaf Khan, Zhengyuan Shi, Naixing Wang, Huang Yu, Qiang Xu</i>	
36.4	Bipolar Vector Classifier for Fault-tolerant Deep Neural Networks	673
	<i>Suyong Lee, Insu Choi, Joon-Sung Yang</i>	
37.1	HDLock: Exploiting Privileged Encoding to Protect Hyperdimensional Computing Models against IP Stealing	679
	<i>Shijin Duan, Shaolei Ren, Xiaolin Xu</i>	
37.2	Terminator on SkyNet: A Practical DVFS Attack on DNN Hardware IP for UAV Object Detection	685
	<i>Junge Xu, Bohan Xuan, Anlin Liu, Mo Sun, Fan Zhang, Zeke Wang, Kui Ren</i>	
37.3	AL-PA: Cross-Device Profiled Side-Channel Attack Using Adversarial Learning	691
	<i>Pei Cao, Hongyi Zhang, Dawu Gu, Yan Lu, Yidong Yuan</i>	
37.4	DETERRENT: Detecting Trojans Using Reinforcement Learning	697
	<i>Vasudev Gohil, Satwik Patnaik, Hao Guo, Dileep Kalathil, Jeyavijayan Rajendran</i>	
38.1	Exploiting Data Locality in Memory for ORAM to Reduce Memory Access Overheads	703
	<i>Jinxi Kuang, Minghua Shen, Yutong Lu, Nong Xiao</i>	
38.2	HWST128: Complete Memory Safety Accelerator on RISC-V with Metadata Compression	709
	<i>Hsu-Kang Dow, Tuo Li, Sri Parameswaran</i>	
38.3	RegVault: Hardware Assisted Selective Data Randomization for Operating System Kernels	715
	<i>Jinyan Xu, Haoran Lin, Ziqi Yuan, Wenbo Shen, Yajin Zhou, Rui Chang, Lei Wu, Kui Ren</i>	
38.4	ASAP: Reconciling Asynchronous Real-Time Operations and Proofs of Execution in Simple Embedded Systems	721
	<i>Adam Caulfield, Norrathep Rattanavipanon, Ivan De Oliveira Nunes</i>	
38.5	Towards a Formally Verified Hardware Root-of-Trust for Data-Oblivious Computing	727
	<i>Lucas Deutschmann, Johannes Müller, Mohammad R. Fadiheh, Dominik Stoffel, Wolfgang Kunz</i>	
38.6	A Scalable SIMD RISC-V based Processor with Customized Vector Extensions for CRYSTALS-Kyber	733
	<i>Huimin Li, Nele Mentens, Stjepan Picek</i>	
39.1	Hexagons are the Bestagons: Design Automation for Silicon Dangling Bond Logic	739
	<i>Marcel Walter, Samuel Sze Hang Ng, Konrad Walus, Robert Wille</i>	
39.2	Improving Compute In-Memory ECC Reliability with Successive Correction	745
	<i>Brian Crafton, Zishen Wan, Samuel Spetalnick, Jong-Hyeok Yoon, Wei Wu, Carlos Tokunaga, Vivek De, Arijit Raychowdhury</i>	

39.3	Energy Efficient Data Search Design and Optimization Based on a Compact Ferroelectric FET Content Addressable Memory	751
	<i>Jiahao Cai, Mohsen Imani, Kai Ni, Grace Li Zhang, Bing Li, Ulf Schlichtmann, Cheng Zhuo, Xunzhao Yin</i>	
39.4	CamSkyGate: Camouflaged Skyrmion Gates for Protecting ICs	757
	<i>Yuqiao Zhang, Chunli Tang, Peng Li, Ujjwal Guin</i>	
39.5	GNN-based Concentration Prediction for Random Microfluidic Mixers	763
	<i>Weiqing Ji, Xingzhuo Guo, Shouan Pan, Tsung-Yi Ho, Ulf Schlichtmann, Hailong Yao</i>	
40.1	Designing ML-Resilient Locking at Register-Transfer Level	769
	<i>Dominik Sisejkovic, Luca Collini, Benjamin Tan, Christian Pilato, Ramesh Karri, Rainer Leupers</i>	
40.2	O'Clock: Lock the Clock via Clock-gating for SoC IP Protection	775
	<i>M. Sazadur Rahman, Rui Guo, Hadi M. Kamali, Fahim Rahman, Farimah Farahmandi, Mohamed Abdel-Moneum, Mark Tehranipoor</i>	
40.3	ALICE: An Automatic Design Flow for eFPGA Redaction	781
	<i>Chiara Muscari Tomajoli, Luca Collini, Jitendra Bhandari, Abdul Khader Thalakkattu Moosa, Benjamin Tan, Xifan Tang, Pierre-Emmanuel Gaillardon, Ramesh Karri, Christian Pilato</i>	
40.4	DELTA: DEsigning a stealthy trigger mechanism for analog hardware Trojans and its detection Analysis	787
	<i>Nishant Gupta, Mohil Sandip Desai, Mark Wijtvliet, Shubham Rai, Akash Kumar</i>	
40.5	VIPR-PCB: A Machine Learning based Golden-Free PCB Assurance Framework	793
	<i>Aritra Bhattacharyay, Prabuddha Chakraborty, Jonathan Cruz, Swarup Bhunia</i>	
40.6	CLIMBER: Defending Phase Change Memory against Inconsistent Write Attacks	799
	<i>Zhuohui Duan, Haobo Wang, Haikun Liu, Xiaofei Liao, Hai Jin, Yu Zhang, Fubing Mao</i>	
41.1	Rethinking Key-Value Store for Byte-Addressable Optane Persistent Memory	805
	<i>Sung-Ming Wu, Li-Pin Chang</i>	
41.2	libcrpm: Improving the Checkpoint Performance of NVM	811
	<i>Feng Ren, Kang Chen, Yongwei Wu</i>	
41.3	Scalable Crash Consistency for Secure Persistent Memory	817
	<i>Ming Zhang, Yu Hua, Xuan Li, Hao Xu</i>	
41.4	Don't Open Row: Rethinking Row Buffer Policy for Improving Performance of Non-Volatile Memories	823
	<i>Yongho Lee, Osang Kwon, Seokin Hong</i>	
41.5	SMART: On Simultaneously MArching RaceTracks to Improve the Performance of RaceTrack-based Main Memory	829
	<i>Xiangjun Peng, Ming-Chang Yang, Ho Ming Tsui, Chi Ngai Leung, Wang Kang</i>	
41.6	SAPredictor: A Simple and Accurate Self-adaptive Predictor for Hierarchical Hybrid Memory System	835
	<i>Yujuan Tan, Wei Chen, Zhulin Ma, Dan Xiao, Zhichao Yan, Duo Liu, Xianzhang Chen</i>	
42.1	AVATAR: An Aging- and Variation-Aware Dynamic Timing Analyzer for Application-based DVAFS	841
	<i>Zuodong Zhang, Zizheng Guo, Yibo Lin, Runsheng Wang, Ru Huang</i>	

42.2	A Defect Tolerance Framework for Improving Yield	847
	<i>Shiva Shankar Thiagarajan, Suriyaprakash Natarajan, Yiorgos Makris</i>	
42.3	Winograd Convolution: A Perspective from Fault Tolerance	853
	<i>Xinghua Xue, Haitong Huang, Cheng Liu, Tao Luo, Lei Zhang, Ying Wang</i>	
42.4	Towards Resilient Analog In-Memory Deep Learning via Data Layout Re-Organization	859
	<i>Muhammad Rashedul Haq Rashed, Amro Awad, Sumit Kumar Jha, Rickard Ewetz</i>	
42.5	SEM-Latch: A Lost-Cost and High-Performance Latch Design for Mitigating Soft Errors in Nanoscale CMOS Process	865
	<i>Zhong-Li Tang, Chia-Wei Liang, Ming-Hsien Hsiao, Charles H.-P. Wen</i>	
43.1	BlueSeer: AI-Driven Environment Detection via BLE Scans	871
	<i>Valentin Poirot, Oliver Harms, Hendric Martens, Olaf Lansiedel</i>	
43.2	Compressive Sensing based Asymmetric Semantic Image Compression for Resource-constrained IoT System	877
	<i>Yujun Huang, Bin Chen, Jianghui Zhang, Qiu Han, Shu-Tao Xia</i>	
43.3	R2B: High-Efficiency and Fair I/O Scheduling for Multi-tenant with Differentiated Demands	883
	<i>Diansen Sun, Yunpeng Chai, Chaoyang Liu, Weihao Sun, Qingpeng Zhang</i>	
43.4	Fast and Scalable Human Pose Estimation Using mmWave Point Cloud	889
	<i>Sizhe An, Umit Y. Ogras</i>	
43.5	VWR2A: A Very-Wide-Register Reconfigurable-Array Architecture for Low-Power Embedded Devices	895
	<i>Benoît W. Denkinger, Miguel Peón-Quirós, Mario Konijnenburg, David Atienza, Francky Catthoor</i>	
43.6	Alleviating Datapath Conflicts and Design Centralization in Graph Analytics Acceleration	901
	<i>Haiyang Lin, Mingyu Yan, Duo Wang, Mo Zou, Fengbin Tu, Xiaochun Ye, Dongrui Fan, Yuan Xie</i>	
44.1	Hyperdimensional Hashing: A Robust and Efficient Dynamic Hash Table	907
	<i>Mike Heddes, Igor Nunes, Tony Givargis, Alexandru Nicolau, Alex Veidenbaum</i>	
44.2	In-Situ Self-Powered Intelligent Vision System with Inference-Adaptive Energy Scheduling for BNN-Based Always-On Perception	913
	<i>Maimaiti Nazhamaiti, Haijin Su, Han Xu, Zheyu Liu, Fei Qiao, Qi Wei, Zidong Du, Xinghua Yang, Li Luo</i>	
44.3	Adaptive Window-Based Sensor Attack Detection for Cyber-Physical Systems	919
	<i>Lin Zhang, Zifan Wang, Mengyu Liu, Fanxin Kong</i>	
44.4	Design-while-Verify: Correct-by-Construction Control Learning with Verification in the Loop	925
	<i>Yixuan Wang, Chao Huang, Zhaoran Wang, Zhilu Wang, Qi Zhu</i>	
45.1	GaBAN: A Generic and Flexibly Programmable Vector Neuro-processor on FPGA	931
	<i>Jiajie Chen, Le Yang, Youhui Zhang</i>	
45.2	ADEPT: Automatic Differentiable DEsign of Photonic Tensor Cores	937
	<i>Jiaqi Gu, Hanqing Zhu, Chenghao Feng, Zixuan Jiang, Mingjie Liu, Shuhan Zhang, Ray T. Chen, David Z. Pan</i>	

45.3	Unicorn: A Multicore Neuromorphic Processor with Flexible Fan-In and Unconstrained Fan-Out for Neurons	943
	<i>Zhijie Yang, Lei Wang, Yao Wang, Linghui Peng, Xiaofan Chen, Xun Xiao, Yaohua Wang, Weixia Xu</i>	
45.4	Effective Zero Compression on ReRAM-based Sparse DNN Accelerators	949
	<i>Hoon Shin, Rihae Park, Seung Yul Lee, Yeonhong Park, Hyunseung Lee, Jae W. Lee</i>	
46.1	Y-architecture-based Flip-Chip Routing with Dynamic Programming-based Bend Minimization	955
	<i>Szu-Ru Nie, Yen-Ting Chen, Yao-Wen Chang</i>	
46.2	Towards Collaborative Intelligence: Routability Estimation based on Decentralized Private Data	961
	<i>Jingyu Pan, Chen-Chia Chang, Zhiyao Xie, Ang Li, Minxue Tang, Tunhou Zhang, Jiang Hu, Yiran Chen</i>	
46.3	A2-ILT: GPU Accelerated ILT with Spatial Attention Mechanism	967
	<i>Qijing Wang, Bentian Jiang, Martin D.F. Wong, Evangeline F.Y. Young</i>	
46.4	Generic Lithography Modeling with Dual-band Optics-Inspired Neural Networks	973
	<i>Haoyu Yang, Zongyi Li, Kumara Sastry, Saumyadip Mukhopadhyay, Mark Kilgard, Anima Anandkumar, Brucek Khailany, Vivek Singh, Haoxing Ren</i>	
47.1	Statistical Computing Framework and Demonstration for In-memory Computing Systems	979
	<i>Bonan Zhang, Peter Deaville, Naveen Verma</i>	
47.2	Write or Not: Programming Scheme Optimization for RRAM-based Neuromorphic Computing	985
	<i>Ziqi Meng, Yanan Sun, Weikang Qian</i>	
47.3	ReSMA: Accelerating Approximate String Matching Using ReRAM-based Content Addressable Memory	991
	<i>Huize Li, Hai Jin, Long Zheng, Yu Huang, Xiaofei Liao, Zhuohui Duan, Dan Chen, Chuangyi Gui</i>	
47.4	VStore: In-Storage Graph Based Vector Search Accelerator	997
	<i>Shengwen Liang, Ying Wang, Ziming Yuan, Cheng Liu, Huawei Li, Xiaowei Li</i>	
48.1	Scaled-CBSC: Scaled Counting-Based Stochastic Computing Multiplication for Improved Accuracy	1003
	<i>Shuyuan Yu, Sheldon X.-D. Tan</i>	
48.2	Tailor: Removing Redundant Operations in Memristive Analog Neural Network Accelerators	1009
	<i>Xingchen Li, Zhihang Yuan, Guangyu Sun, Liang Zhao, Zhichao Lu</i>	
48.3	Domain Knowledge-Infused Deep Learning for Automated Analog/Radio-Frequency Circuit Parameter Optimization	1015
	<i>Weidong Cao, Mouhacine Benosman, Xuan Zhang, Rui Ma</i>	
48.4	A Cost-Efficient Fully Synthesizable Stochastic Time-to-Digital Converter Design Based on Integral Nonlinearity Scrambling	1021
	<i>Qiaochu Zhang, Shiyu Su, Mike Shuo-Wei Chen</i>	
49.1	Using Machine Learning to Optimize Graph Execution on NUMA Machines	1027
	<i>Hiago Mayk G. de A. Rocha, Janaina Schwarzrock, Arthur F. Lorenzon, Antonio Carlos Schneider Beck</i>	
49.2	HCG: Optimizing Embedded Code Generation of Simulink with SIMD Instruction Synthesis	1033
	<i>Zhuo Su, Zehong Yu, Dongyan Wang, Yixiao Yang, Yu Jiang, Rui Wang, Wanli Chang, Jiaguang Sun</i>	

49.3	Raven: A Novel Kernel Debugging Tool on RISC-V	1039
	<i>Hongyi Lu, Fengwei Zhang</i>	
49.4	GTuner: Tuning DNN Computations on GPU via Graph Attention Network	1045
	<i>Qi Sun, Xinyun Zhang, Hao Geng, Yuxuan Zhao, Yang Bai, Haisheng Zheng, Bei Yu</i>	
50.1	Pref-X: A Framework to Reveal Data Prefetching in Commercial In-Order Cores	1051
	<i>Quentin Huppert, Francky Catthoor, Lionel Torres, David Novo</i>	
50.2	Architecting DDR5 DRAM Caches for Non-Volatile Memory Systems	1057
	<i>Xin Xin, Wanyi Zhu, Li Zhao</i>	
50.3	GraphRing: An HMC-Ring based Graph Processing Framework with Optimized Data Movement	1063
	<i>Zerun Li, Xiaoming Chen, Yinhe Han</i>	
50.4	AxoNN: Energy-Aware Execution of Neural Network Inference on Multi-Accelerator Heterogeneous SoCs	1069
	<i>Ismet Dagli, Alexander Cieslewicz, Jedidiah McClurg, Mehmet E. Belviranli</i>	
51.1	PIPF-DRAM: Processing in Precharge-Free DRAM	1075
	<i>Nezam Rohbani, Mohammad Arman Soleimani, Hamid Sarbazi-Azad</i>	
51.2	TAIM: Ternary Activation In-Memory Computing Hardware with 6T SRAM Array	1081
	<i>Nameun Kang, Hyungjun Kim, Hyunmyung Oh, Jae-Joon Kim</i>	
51.3	PIM-DH: ReRAM-based Processing-in-Memory Architecture for Deep Hashing Acceleration	1087
	<i>Fangxin Liu, Wenbo Zhao, Yongbiao Chen, Zongwu Wang, Zhezhi He, Rui Yang, Qidong Tang, Tao Yang, Cheng Zhuo, Li Jiang</i>	
51.4	YOLoC: DeploY Large-Scale Neural Network by ROM-based Computing-in-Memory using Residual Branch on a Chip	1093
	<i>Yiming Chen, Guodong Yin, Zhanhong Tan, Mingyen Lee, Zekun Yang, Yongpan Liu, Huazhong Yang, Kaisheng Ma, Xueqing Li</i>	
52.1	ASTERS: Adaptable Threshold Spike-timing Neuromorphic Design with Twin-Column ReRAM Synapses ...	1099
	<i>Ziru Li, Qilin Zheng, Bonan Yan, Ru Huang, Bing Li, Yiran Chen</i>	
52.2	SATO: Spiking Neural Network Acceleration via Temporal-Oriented Dataflow and Architecture	1105
	<i>Fangxin Liu, Wenbo Zhao, Zongwu Wang, Yongbiao Chen, Tao Yang, Zhezhi He, Xiaokang Yang, Li Jiang</i>	
52.3	LeHDC: Learning-Based Hyperdimensional Computing Classifier	1111
	<i>Shijin Duan, Yejia Liu, Shaolei Ren, Xiaolin Xu</i>	
52.4	GENERIC: Highly Efficient Learning Engine on Edge Using Hyperdimensional Computing	1117
	<i>Behnam Khaleghi, Jaeyoung Kang, Hanyang Xu, Justin Morris, Tajana Rosing</i>	
52.5	Solving Traveling Salesman Problems via a Parallel Fully Connected Ising Machine	1123
	<i>Qichao Tao, Jie Han</i>	
52.6	PATH: Evaluation of Boolean Logic Using Path-based In-Memory Computing	1129
	<i>Sven Thijssen, Sumit Kumar Jha, Rickard Ewetz</i>	

53.1	A Length Adaptive Algorithm-Hardware Co-design of Transformer on FPGA through Sparse Attention and Dynamic Pipelining	1135
	<i>Hongwu Peng, Shaoyi Huang, Shiyang Chen, Bingbing Li, Tong Geng, Ang Li, Weiwen Jiang, Wujie Wen, Jinbo Bi, Hang Liu, Caiwen Ding</i>	
53.2	HDPG: Hyperdimensional Policy-based Reinforcement Learning for Continuous Control	1141
	<i>Yang Ni, Mariam Issa, Danny Abraham, Mahdi Imani, Xunzhao Yin, Mohsen Imani</i>	
53.3	CarM: Hierarchical Episodic Memory for Continual Learning	1147
	<i>Soobee Lee, Minindu Weerakoon, Jonghyun Choi, Minjia Zhang, Di Wang, Myeongjae Jeon</i>	
53.4	Shfl-BW: Accelerating Deep Neural Network Inference with Tensor-Core Aware Weight Pruning	1153
	<i>Guyue Huang, Haoran Li, Minghai Qin, Fei Sun, Yufei Ding, Yuan Xie</i>	
53.5	QuiltNet: Efficient Deep Learning Inference on Multi-Chip Accelerators Using Model Partitioning	1159
	<i>Jongho Park, HyukJun Kwon, Seowoo Kim, Junyoung Lee, Minh Ho, Euicheol Lim, Mohsen Imani, Yeseong Kim</i>	
53.6	Glimpse: Mathematical Embedding of Hardware Specification for Neural Compilation	1165
	<i>Byung Hoon Ahn, Sean Kinzer, Hadi Esmaeilzadeh</i>	
54.1	Bringing Source-Level Debugging Frameworks to Hardware Generators	1171
	<i>Keyi Zhang, Zain Asgar, Mark Horowitz</i>	
54.2	Verifying SystemC TLM Peripherals Using Modern C++ Symbolic Execution Tools	1177
	<i>Pascal Pieper, Vladimir Herdt, Daniel Große, Rolf Drechsler</i>	
54.3	Formal Verification of Modular Multipliers Using Symbolic Computer Algebra and Boolean Satisfiability	1183
	<i>Alireza Mahzoon, Daniel Große, Christoph Scholl, Alexander Konrad, Rolf Drechsler</i>	
54.4	Silicon Validation of LUT-based Logic-Locked IP Cores	1189
	<i>Gaurav Kolhe, Tyler Sheaves, Kevin Immanuel Gubbi, Tejas Kadale, Setareh Rafatirad, Sai Manoj Pudukotai Dinakarrao, Avesta Sasan, Hamid Mahmoodi, Houman Homayoun</i>	
54.5	Efficient Bayesian Yield Analysis and Optimization with Active Learning	1195
	<i>Shuo Yin, Xiang Jin, Linxu Shi, Kang Wang, Wei W. Xing</i>	
54.6	Accelerated Synthesis of Neural Network-based Barrier Certificates Using Collaborative Learning	1201
	<i>Jun Xia, Ming Hu, Xin Chen, Mingsong Chen</i>	
55.1	A Timing Engine Inspired Graph Neural Network Model for Pre-Routing Slack Prediction	1207
	<i>Zizheng Guo, Mingjie Liu, Jiaqi Gu, Shuhan Zhang, David Z. Pan, Yibo Lin</i>	
55.2	Accurate Timing Prediction at Placement Stage with Look-Ahead RC Network	1213
	<i>Xu He, Zhiyong Fu, Yao Wang, Chang Liu, Yang Guo</i>	
55.3	Timing Macro Modeling with Graph Neural Networks	1219
	<i>Kevin Kai-Chun Chang, Chun-Yao Chiang, Pei-Yu Lee, Iris Hui-Ru Jiang</i>	
55.4	Worst-Case Dynamic Power Distribution Network Noise Prediction Using Convolutional Neural Network	1225
	<i>Xiao Dong, Yufei Chen, Xunzhao Yin, Cheng Zhuo</i>	

55.5	GATSPI: GPU Accelerated Gate-Level Simulation for Power Improvement	1231
	<i>Yanqing Zhang, Haoxing Ren, Akshay Sridharan, Brucek Khailany</i>	
55.6	PPATuner: Pareto-driven Tool Parameter Auto-tuning in Physical Design via Gaussian Process Transfer Learning	1237
	<i>Hao Geng, Qi Xu, Tsung-Yi Ho, Bei Yu</i>	
56.1	Efficient Maximum Data Age Analysis for Cause-Effect Chains in Automotive Systems	1243
	<i>Ran Bi, Xinbin Liu, Jiankang Ren, Pengfei Wang, Huawei Lv, Guozhen Tan</i>	
56.2	Optimizing Parallel PREM Compilation over Nested Loop Structures	1249
	<i>Zhao Gu, Rodolfo Pellizzoni</i>	
56.3	Scheduling and Analysis of Real-Time Tasks with Parallel Critical Sections	1255
	<i>Yang Wang, Xu Jiang, Nan Guan, Mingsong Lv, Dong Ji, Wang Yi</i>	
56.4	BlueScale: A Scalable Memory Architecture for Predictable Real-Time Computing on Highly Integrated SoCs	1261
	<i>Zhe Jiang, Kecheng Yang, Neil Audsley, Nathan Fisher, Weisong Shi, Zheng Dong</i>	
56.5	Precise and Scalable Shared Cache Contention Analysis for WCET Estimation	1267
	<i>Wei Zhang, Mingsong Lv, Wanli Chang, Lei Ju</i>	
56.6	Predictable Sharing of Last-level Cache Partitions for Multi-core Safety-critical Systems	1273
	<i>Zhuanhao Wu, Hiren Patel</i>	
57.1	Thermal-aware Optical-electrical Routing Codesign for On-Chip Signal Communications	1279
	<i>Yu-Sheng Lu, Kuan-Cheng Chen, Yu-Ling Hsu, Yao-Wen Chang</i>	
57.2	Power-Aware Pruning for Ultrafast, Energy-Efficient, and Accurate Optical Neural Network Design	1285
	<i>Naoki Hattori, Yutaka Masuda, Tohru Ishihara, Akihiko Shinya, Masaya Notomi</i>	
57.3	REACT: A Heterogeneous Reconfigurable Neural Network Accelerator with Software-Configurable NoCs for Training and Inference on Wearables	1291
	<i>Mohit Upadhyay, Rohan Juneja, Bo Wang, Jun Zhou, Weng-Fai Wong, Li-Shiuan Peh</i>	
57.4	LHNN: Lattice Hypergraph Neural Network for VLSI Congestion Prediction	1297
	<i>Bowen Wang, Guibao Shen, Dong Li, Jianye Hao, Wulong Liu, Yu Huang, Hongzhong Wu, Yibo Lin, Guangyong Chen, Pheng Ann Heng</i>	
58.1	Floorplanning with Graph Attention	1303
	<i>Yiting Liu, Ziyi Ju, Zhengming Li, Mingzhi Dong, Hai Zhou, Jia Wang, Fan Yang, Xuan Zeng, Li Shang</i>	
58.2	Xplace: An Extremely Fast and Extensible Global Placement Framework	1309
	<i>Lixin Liu, Bangqi Fu, Martin D.F. Wong, Evangeline F.Y. Young</i>	
58.3	Differentiable-Timing-Driven Global Placement	1315
	<i>Zizheng Guo, Yibo Lin</i>	
58.4	TAAS: A Timing-Aware Analytical Strategy for AQFP-Capable Placement Automation	1321
	<i>Peiyan Dong, Yanyue Xie, Hongjia Li, Mengshu Sun, Olivia Chen, Nobuyuki Yoshikawa, Yanzhi Wang</i>	

59.1	Invited: A Cross-layer Approach to Cognitive Computing	1327
	<i>Gobinda Saha, Cheng Wang, Anand Raghunathan, Kaushik Roy</i>	
60.2	Invited: Generative Self-Supervised Learning for Gate Sizing	1331
	<i>Siddhartha Nath, Geraldo Pradipta, Corey Hu, Tian Yang, Brucek Khailany, Haoxing Ren</i>	
62.1	Invited: Hammer: A Modular and Reusable Physical Design Flow Tool	1335
	<i>Harrison Liew, Daniel Grubb, John Wright, Colin Schmidt, Nayiri Krzysztofowicz, Adam Izraelevitz, Edward Wang, Krste Asanović, Jonathan Bachrach, Borivoje Nikolić</i>	
62.2	Invited: mflowgen: A Modular Flow Generator and Ecosystem for Community-Driven Physical Design	1339
	<i>Alex Carsello, James Thomas, Ankita Nayak, Po-Han Chen, Mark Horowitz, Priyanka Raina, Christopher Torng</i>	
62.3	Invited: A Distributed Approach to Silicon Compilation	1343
	<i>Andreas Olofsson, William Ransohoff, Noah Moroze</i>	
63.1	Invited: Improving GNN-Based Accelerator Design Automation with Meta Learning	1347
	<i>Yunsheng Bai, Atefeh Sohrabizadeh, Yizhou Sun, Jason Cong</i>	
63.2	Invited: Accelerator Design with Decoupled Hardware Customizations: Benefits and Challenges	1351
	<i>Debjit Pal, Yi-Hsiang Lai, Shaojie Xiang, Niansong Zhang, Hongzheng Chen, Jeremy Casas, Pasquale Cocchini, Zhenkun Yang, Jin Yang, Louis-Noël Pouchet, Zhiru Zhang</i>	
63.3	Invited: ScaleHLS: A Scalable High-Level Synthesis Framework with Multi-level Transformations and Optimizations	1355
	<i>Hanchen Ye, HyeGang Jun, Hyunmin Jeong, Stephen Neuendorffer, Deming Chen</i>	
63.4	Invited: The SODA Approach: Leveraging High-Level Synthesis for Hardware/Software Co-design and Hardware Specialization	1359
	<i>Nicolas Bohm Agostini, Serena Curzel, Ankur Limaye, Vinay Amatya, Marco Minutoli, Vito Giovanni Castellana, Joseph Manzano, Antonino Tumeo, Fabrizio Ferrandi</i>	
64.1	Invited: Automatic Oracle Generation in Microsoft's Quantum Development Kit Using QIR and LLVM Passes	1363
	<i>Mathias Soeken, Mariia Mykhailova</i>	
64.2	Invited: The Basis of Design Tools for Quantum Computing: Arrays, Decision Diagrams, Tensor Networks, and ZX-Calculus	1367
	<i>Robert Wille, Lukas Burgholzer, Stefan Hillmich, Thomas Grurl, Alexander Ploier, Tom Peham</i>	
66.1	Invited: Secure by Construction: Addressing Security Vulnerabilities Introduced During High-level Synthesis	1371
	<i>Md Rafid Muttaki, Zahin Ibnat, Farimah Farahmandi</i>	
66.2	Invited: High-Level Design Methods for Hardware Security: Is It the Right Choice?	1375
	<i>Christian Pilato, Donatella Sciuto, Benjamin Tan, Siddharth Garg, Ramesh Karri</i>	
66.3	Invited: Trusting the Trust Anchor: Towards Detecting Cross-Layer Vulnerabilities with Hardware Fuzzing	1379
	<i>Chen Chen, Rahul Kande, Pouya Mahmoody, Ahmad-Reza Sadeghi, JV Rajendran</i>	
66.4	Invited: Automating Hardware Security Property Generation	1384
	<i>Ryan Kastner, Francesco Restuccia, Andres Meza, Sayak Ray, Jason Fung, Cynthia Sturton</i>	

67.1	Late Breaking Results: Efficient Timing Propagation with Simultaneous Structural and Pipeline Parallelisms	1388
	<i>Cheng-Hsiang Chiu, Tsung-Wei Huang</i>	
67.2	Late Breaking Results: A Fast and Low-Cost Comparison-Free Sorting Engine with Unary Computing	1390
	<i>Amir Hossein Jalilvand, Seyedeh Newsha Estiri, Samaneh Naderi, M. Hassan Najafi, Mohsen Imani</i>	
67.3	Late Breaking Results: Flexible Chip Placement via Reinforcement Learning	1392
	<i>Fu-Chieh Chang, Yu-Wei Tseng, Ya-Wen Yu, Ssu-Rui Lee, Alexandra Cioba, I-Lun Tseng, Da-Shan Shiu, Jhih-Wei Hsu, Cheng-Yuan Wang, Chien-Yi Yang, Ren-Chu Wang, Yao-Wen Chang, Tai-Chen Chen, Tung-Chieh Chen</i>	
67.4	Late Breaking Results: FPGA-Aware Automatic Acceleration Framework for Vision Transformer with Mixed-Scheme Quantization	1394
	<i>Mengshu Sun, Zhengang Li, Alec Lu, Haoyu Ma, Geng Yuan, Yanyue Xie, Hao Tang, Yanyu Li, Miriam Leiser, Zhangyang Wang, Xue Lin, Zhenman Fang</i>	
67.5	Late Breaking Results: Hardware-Efficient Stochastic Rounding Unit Design for DNN Training	1396
	<i>Sung-En Chang, Geng Yuan, Alec Lu, Mengshu Sun, Yanyu Li, Xiaolong Ma, Zhengang Li, Yanyue Xie, Minghai Qin, Xue Lin, Zhenman Fang, Yanzhi Wang</i>	
67.6	Late Breaking Results: Placement Initialization via a Projected Eigenvector Algorithm	1398
	<i>Pengwen Chen, Chung-Kuan Cheng, Albert Chern, Chester Holtz, Aoxi Li, Yucheng Wang</i>	
67.7	Late Breaking Results: Subgraph Matching Based Reference Placement for PCB Designs	1400
	<i>Miaodi Su, Yifeng Xiao, Shu Zhang, Haiyuan Su, Jiachen Xu, Huan He, Ziran Zhu, Jianli Chen, Yao-Wen Chang</i>	
67.8	Late Breaking Results: Thermal-Aware Drone Battery Management	1402
	<i>Hojun Choi, Youngmoon Lee</i>	
67.9	Late Breaking Results: Waveform-based Performance Analysis of RISC-V Processors	1404
	<i>Lucas Klemmer, Daniel Große</i>	
	ADA, the Center for Applications Driving Architectures: Accomplishments and Vision Forward	1406
	<i>Valeria Bertacco, Todd Austin, David Brooks, Sharad Malik, Zachary Tatlock, Gu-Yeon Wei, Thomas F. Wenisch</i>	