

# **2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC 2024)**

**Incheon, South Korea  
22-25 January 2024**

**Pages 1-490**



**IEEE Catalog Number: CFP24ASP-POD  
ISBN: 979-8-3503-9355-2**

**Copyright © 2024 by the Institute of Electrical and Electronics Engineers, Inc.  
All Rights Reserved**

*Copyright and Reprint Permissions:* Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limit of U.S. copyright law for private use of patrons those articles in this volume that carry a code at the bottom of the first page, provided the per-copy fee indicated in the code is paid through Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

For other copying, reprint or republication permission, write to IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, Piscataway, NJ 08854. All rights reserved.

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

IEEE Catalog Number:	CFP24ASP-POD
ISBN (Print-On-Demand):	979-8-3503-9355-2
ISBN (Online):	979-8-3503-9354-5
ISSN:	2153-6961

**Additional Copies of This Publication Are Available From:**

Curran Associates, Inc  
57 Morehouse Lane  
Red Hook, NY 12571 USA  
Phone: (845) 758-0400  
Fax: (845) 758-2633  
E-mail: [curran@proceedings.com](mailto:curran@proceedings.com)  
Web: [www.proceedings.com](http://www.proceedings.com)

CURRAN ASSOCIATES INC.  
**proceedings**  
.com

# ASP-DAC 2024

## Contents

### Front Matter

Welcome to ASP-DAC 2024 . . . . .	viii
Message from the Technical Program Committee . . . . .	ix
Organizing Committee . . . . .	x
Technical Program Committee . . . . .	xi
University LSI Design Contest Committee . . . . .	xiii
Designers' Forum Committee . . . . .	xiv
Steering Committee . . . . .	xv
University LSI Design Contest . . . . .	xvi
Designers' Forum . . . . .	xvii
ACM SIGDA Student Research Forum at ASP-DAC 2024 . . . . .	xix
Best Paper Award . . . . .	xx
University LSI Design Contest Award . . . . .	xxi
10-Year Retrospective Most Influential Paper Award . . . . .	xxii
Keynote Addresses . . . . .	xxiii
Invitation to ASP-DAC 2025 . . . . .	xxv
List of Reviewers (Regular topic) . . . . .	xxvi
List of Reviewers (UDC) . . . . .	xxvii

### Technical Program

#### Session 1A: Emerging NoC Designs

CANSim: When to Utilize Synchronous and Asynchronous Routers in Large and Complex NoCs . . . . .	1
PAIR: Periodically Alternate the Identity of Routers to Ensure Deadlock Freedom in NoC . . . . .	7
SCNoCs: An Adaptive Heterogeneous Multi-NoC with Selective Compression and Power Gating . . . . .	13

#### Session 1B: When AI Meets Edge Devices

QuadraNet: Improving High-Order Neural Interaction Efficiency with Hardware-Aware Quadratic Neural Networks . . . . .	19
RobustDiCE: Robust and Distributed CNN Inference at the Edge . . . . .	26
YoseUe: "trimming" Random Forest's training towards resource-constrained inference . . . . .	32

#### Session 1C: Innovations in New Computing Paradigms: Stochastic, Hyper-Dimensional and High-Performance Computing

P2LSG: Powers-of-2 Low-Discrepancy Sequence Generator for Stochastic Computing . . . . .	38
--	----

PAAP-HD: PIM-Assisted Approximation for Efficient Hyper-Dimensional Computing . . . . .	46
FPGA-Based HPC for Associative Memory System . . . . .	52
<b>Session 1D: 3D IC</b>	
Chiptetizer: Repartitioning SoCs for Cost-Effective Chiplet Integration . . . . .	58
CoPlace: Coherent Placement Engine with Layout-aware Partitioning for 3D ICs . . . . .	65
O.O: Optimized One-die Placement for Face-to-face Bonded 3D ICs . . . . .	71
<b>Session 1E: Open-Source EDA Algorithms and Software</b>	
iEDA: An Open-source infrastructure of EDA . . . . .	77
iPD: An Open-source intelligent Physical Design Toolchain . . . . .	83
A Resource-efficient Task Scheduling System Using Reinforcement Learning . . . . .	89
Machine learning and GPU accelerated sparse linear solvers for transistor-level circuit simulation: a perspective survey . . . . .	96
<b>Session 2A: Frontiers in Embedded and Edge AI: from Adversarial Attacks to Intelligent Homes</b>	
HomeSGN: A Smarter Home with Novel Rule Mining Enabled by a Scorer-Generator GAN . . . . .	102
Adaptive Workload Distribution for Accuracy-aware DNN Inference on Collaborative Edge Platforms . . . . .	109
Extending Neural Processing Unit and Compiler for Advanced Binarized Neural Networks . . . . .	115
<b>Session 2B: Innovations in Quantum EDA: from Design to Deployment</b>	
JustQ: Automated Deployment of Fair and Accurate Quantum Neural Networks . . . . .	121
Using Boolean Satisfiability for Exact Shuttling in Trapped-Ion Quantum Computers . . . . .	127
Optimizing Decision Diagrams for Measurements of Quantum Circuits . . . . .	134
CTQr: Control and Timing-Aware Qubit Routing . . . . .	140
<b>Session 2C: Architecting for Dependability: System Design with Compute-in-Memory</b>	
BNN-Flip: Enhancing the Fault Tolerance and Security of Compute-in-Memory Enabled Binary Neural Network Accelerators . . . . .	146
ZEBRA: A Zero-Bit Robust-Accumulation Compute-In-Memory Approach for Neural Network Acceleration Utilizing Different Bitwise Patterns . . . . .	153
A Cross-layer Framework for Design Space and Variation Analysis of Non-Volatile Ferroelectric Capacitor-Based Compute-in-Memory Accelerators . . . . .	159
Design of Aging-Robust Clonable PUF Using an Insulator-Based ReRAM for Organic Circuits . . . . .	165
<b>Session 2D: ML for Physical Design and Timing</b>	
Heterogeneous Graph Attention Network Based Statistical Timing Library Characterization with Parasitic RC Reduction . . . . .	171
An Optimization-aware Pre-Routing Timing Prediction Framework Based on Heterogeneous Graph Learning . . . . .	177
BoCNT: A Bayesian Optimization Framework for Global CNT Interconnect Optimization . . . . .	183
Timing Analysis beyond Complementary CMOS Logic Styles . . . . .	189
<b>Session 3A: GPU and Custom Accelerators</b>	

Collaborative Coalescing of Redundant Memory Access for GPU System . . . . .	195
WER: Maximizing Parallelism of Irregular Graph Applications Through GPU Warp EqualizeR . . . . .	201
SoC-Tuner: An Importance-guided Exploration Framework for DNN-targeting SoC Design . . . . .	207
ARS-Flow: A Design Space Exploration Flow for Accelerator-rich System based on Active Learning . . . . .	213
Secco: Codesign for Resource Sharing in Regular-Expression Accelerators . . . . .	219
<b>Session 3B: Hardware Acceleration for Graph Neural Networks and New Models</b>	
SparGNN: Efficient Joint Feature-Model Sparsity Exploitation in Graph Neural Network Acceleration . . . . .	225
APoX: Accelerate Graph-Based Deep Point Cloud Analysis via Adaptive Graph Construction . . . . .	231
FuseFPS: Accelerating Farthest Point Sampling with Fusing KD-tree Construction for Point Clouds . . . . .	238
A Fixed-Point Pre-Processing Hardware Architecture Design for Complex Independent Component Analysis . . . . .	244
Pearls Hide Behind Linearity: Simplifying Deep Convolutional Networks for Embedded Hardware Systems via Linearity Grafting . . . . .	250
<b>Session 3C: New Frontiers in Verification and Simulation</b>	
On Decomposing Complex Test Cases for Efficient Post-silicon Validation . . . . .	256
DeepIC3: Guiding IC3 Algorithms by Graph Neural Network Clause Prediction . . . . .	262
TIUP: Effective Processor Verification with Tautology-Induced Universal Properties . . . . .	269
Verifying Embedded Graphics Libraries leveraging Virtual Prototypes and Metamorphic Testing . . . . .	275
MemSPICE: Automated Simulation and Energy Estimation Framework for MAGIC-Based Logic-in-Memory . . . . .	282
<b>Session 3D: Partition and Placement</b>	
An Effective Netlist Planning Approach for Double-sided Signal Routing . . . . .	288
An Analytical Placement Algorithm with Routing topology Optimization . . . . .	294
Effective Analytical Placement for Advanced Hybrid-Row-Height Circuit Designs . . . . .	300
Row Planning and Placement for Hybrid-Row-Height Designs . . . . .	306
TransPlace: A Scalable Transistor-Level Placer for VLSI Beyond Standard-Cell-Based Design . . . . .	312
<b>Session 3E: Hardware Security – A True Multidisciplinary Research Area</b>	
Towards Finding the Sources of Polymorphism in Polymorphic Gates . . . . .	319
HOGGE: Homomorphic Gate on An FPGA . . . . .	325
Sensors for Remote Power Attacks: New Developments and Challenges . . . . .	333
PRESS: Persistence Relaxation for Efficient and Secure Data Sanitization on Zoned Namespace Storage . . . . .	341
Hardware Phi-1.5B: A Large Language Model Encodes Hardware Domain Specific Knowledge . . . . .	349
<b>Session 4A: Detection Techniques for SoC Vulnerability and Malware</b>	
FormalFuzzer: Formal Verification Assisted Fuzz Testing for SoC Vulnerability Detection . . . . .	355
DeepIncept: Diversify Performance Counters with Deep Learning to Detect Malware . . . . .	362
Resource- and Workload-aware Malware Detection through Distributed Computing in IoT Networks . . . . .	368

<b>Session 4B: Design for Manufacturability: from Rule Checking to Yield Optimization</b>	
APPLE: An Explainer of ML Predictions on Circuit Layout at the Circuit-Element Level . . . . .	374
E2E-Check: End to End GPU-Accelerated Design Rule Checking with Novel Mask Boolean Algorithms	380
CIS: Conditional Importance Sampling for Yield Optimization of Analog and SRAM Circuits . . . . .	386
<b>Session 4C: Advances in Logic Synthesis and Optimization</b>	
FineMap: A Fine-grained GPU-parallel LUT Mapping Engine . . . . .	392
Transduction Method for AIG Minimization . . . . .	398
In Medio Stat Virtus: Combining Boolean and Pattern Matching . . . . .	404
<b>Session 4D: Learning-Based Optimization for RF/Analog Circuit</b>	
A Transferable GNN-based Multi-Corner Performance Variability Modeling for Analog ICs . . . . .	411
An Efficient Transfer Learning Assisted Global Optimization Scheme for Analog/RF Circuits . . . . .	417
MACRO: Multi-agent Reinforcement Learning-based Cross-layer Optimization of Operational Amplifier	423
<b>Session 4E: LLM Acceleration and Specialization for Circuit Design and Edge Applications</b>	
FL-NAS: Towards Fairness of NAS for Resource Constrained Devices via Large Language Models . . . . .	429
Software/Hardware Co-design for LLM and Its Application for Design Verification . . . . .	435
<b>Session 5A: Bridging Memory, Storage, and Data Processing Techniques</b>	
wearMeter: an Accurate Wear Metric for NAND Flash Memory . . . . .	442
Overlapping Aware Zone Allocation for LSM Tree-Based Store on ZNS SSDs . . . . .	448
Hardware-Software Co-Design of a Collaborative DNN Accelerator for 3D Stacked Memories with Multi-Channel Data . . . . .	454
Bridge-NDP: Achieving Efficient Communication-Computation Overlap in Near Data Processing with Bridge Architecture . . . . .	460
<b>Session 5B: Exploring EDA’s Next Frontier: AI-Driven Innovative Design Methods</b>	
Variational Label-Correlation Enhancement for Congestion Prediction . . . . .	466
Fast Cell Library Characterization for Design Technology Co-Optimization Based on Graph Neural Networks . . . . .	472
Automated synthesis of mixed-signal ML inference hardware under accuracy constraints . . . . .	478
LayNet: Layout Size Prediction for Memory Design Using Graph Neural Networks in Early Design Stage . . . . .	484
<b>Session 5C: New Frontiers in Testing</b>	
QcAssert: Quantum Device Testing with Concurrent Assertions . . . . .	491
HybMT: Hybrid Meta-Predictor based ML Algorithm for Fast Test Vector Generation . . . . .	497
A Fast Test Compaction Method for Commercial DFT Flow Using Dedicated Pure-MaxSAT Solver . . . . .	503
A Dynamic Testing Scheme for Resistive-Based Computation-In-Memory Architectures . . . . .	509
<b>Session 5D: FPGA-Based Neural Network Accelerator Designs and Applications</b>	
SWAT: An Efficient Swin Transformer Accelerator Based on FPGA . . . . .	515

TransFRU: Efficient Deployment of Transformers on FPGA with Full Resource Utilization . . . . .	521
Booth-NeRF: An FPGA Accelerator for Instant-NGP Inference with Novel Booth-Multiplier . . . . .	527
ACane: An Efficient FPGA-based Embedded Vision Platform with Accumulation-as-Convolution Packing for Autonomous Mobile Robots . . . . .	533
<b>Session 6A: Enabling Techniques to Make CIM Small and Flexible</b>	
OSA-HCIM: On-The-Fly Saliency-Aware Hybrid SRAM CIM with Dynamic Precision Configuration	539
BFP-CIM: Data-Free Quantization with Dynamic Block-Floating-Point Arithmetic for Energy- Efficient Computing-In-Memory-based Accelerator . . . . .	545
A Heuristic and Greedy Weight Remapping Scheme with Hardware Optimization for Irregular Sparse Neural Networks Implemented on CIM Accelerator in Edge AI Applications . . . . .	551
PRIMATE: Processing in Memory Acceleration for Dynamic Token-pruning Transformers . . . . .	557
<b>Session 6B: Emerging Trends in Hardware Design: from Biochips to Quantum Systems</b>	
Adaptive Control-Logic Routing for Fully Programmable Valve Array Biochips Using Deep Rein- forcement Learning . . . . .	564
Towards Automated Testing of Multiplexers in Fully Programmable Valve Array Biochips . . . . .	570
The Need for Speed: Efficient Exact Simulation of Silicon Dangling Bond Logic . . . . .	576
Towards Multiphase Clocking in Single-Flux Quantum Systems . . . . .	582
Algebraic and Boolean Methods for SFQ Superconducting Circuits . . . . .	588
<b>Session 6C: New Advances in Logic Locking and Side-Channel Analysis</b>	
LOOPLock 3.0: A Robust Cyclic Logic Locking Approach . . . . .	594
Logic Locking over TFHE for Securing User Data and Algorithms . . . . .	600
LIPSTICK: Corruptibility-Aware and Explainable Graph Neural Network-based Oracle-Less Attack on Logic Locking . . . . .	606
Power Side-Channel Analysis and Mitigation for Neural Network Accelerators based on Memristive Crossbars . . . . .	612
Modeling of Tamper Resistance to Electromagnetic Side-channel Attacks on Voltage-scaled Circuits .	618
<b>Session 6D: Advanced Simulation and Modeling</b>	
SPIRAL: Signal-Power Integrity Co-Analysis for High-Speed Inter-Chiplet Serial Links Validation . .	625
Nested Dissection Based Parallel Transient Power Grid Analysis on Public Cloud Virtual Machines .	631
Efficient Sublogic-Cone-Based Switching Activity Estimation using Correlation Factor . . . . .	638
ISOP-Yield: Yield-Aware Stack-Up Optimization for Advanced Package using Machine Learning . .	644
Physics-Informed Learning for EPG-Based TDDDB Assessment . . . . .	651
<b>Session 6E: Cutting-Edge Techniques for EDA in Analog/Mixed-Signal ICs</b>	
Toward End-to-End Analog Design Automation with ML and Data-Driven Approaches . . . . .	657
Reinforcing the Connection between Analog Design and EDA . . . . .	665
A Study on Exploring and Exploiting the High-dimensional Design Space for Analog Circuit Design Automation . . . . .	671

Performance-Driven Analog Layout Automation: Current Status and Future Directions . . . . .	679
<b>Session 7A: System Performance and Debugging</b>	
The Optimal Quantum of Temporal Decoupling . . . . .	686
Towards a Highly Interactive Design-Debug-Verification Cycle . . . . .	692
Beyond Time-Quantum: A Basic-Block FDA Approach for Accurate System Computing Performance Estimation . . . . .	698
<b>Session 7B: Innovations in Autonomous Systems: Hyperdimensional Computing and Emerging Application Frontiers</b>	
BoostIID: Fault-agnostic Online Detection of WCET Changes in Autonomous Driving . . . . .	704
KalmanHD: Robust On-Device Time Series Forecasting with Hyperdimensional Computing . . . . .	710
HyperFeel: An Efficient Federated Learning Framework Using Hyperdimensional Computing . . . . .	716
<b>Session 7C: Productivity Management for High Level Design</b>	
RTLML: An Open-Source Benchmark for Design RTL Generation with Large Language Model . . . . .	722
LSTP : A Logic Synthesis Timing Predictor . . . . .	728
Bridging the Design Methodologies of Burst-Mode Specifications and Signal Transition Graphs . . . . .	734
<b>Session 7D: Emerging Memory Yield Optimization and Modeling</b>	
Signature Driven Post-Manufacture Testing and Tuning of RRAM Spiking Neural Networks for Yield Recovery . . . . .	740
Physics-Informed Learning for Versatile RRAM Reset and Retention Simulation . . . . .	746
Hard Error Correction in STT-MRAM . . . . .	752
<b>Session 7E: Enabling Chiplet-based Custom Designs</b>	
Exploiting 2.5D/3D Heterogeneous Integration for AI Computing . . . . .	758
Challenges and Opportunities to Enable Large-scale Computing via Heterogeneous Chiplets . . . . .	765
Towards Automated Generation of Chiplet-Based System . . . . .	771
<b>Session 8A: Advances in Efficient Embedded Computing: from Hardware Accelerator to Task Management</b>	
Flexible Spatio-Temporal Energy-Efficient Runtime Management . . . . .	777
Sparse-Sparse Matrix Multiplication Accelerator on FPGA featuring Distribute-Merge Product Dataflow	785
Meeting Job-Level Dependencies by Task Merging . . . . .	792
A CGRA Front-end Compiler Enabling Extraction of General Control and Dedicated Operators . . . . .	799
<b>Session 8B: In-Memory Computing Architecture Design and Logic Synthesis</b>	
LOSSS- Logic Synthesis based on Several Stateful logic gates for high time-efficient computing . . . . .	805
Towards Area-Efficient Path-Based In-Memory Computing using Graph Isomorphisms . . . . .	812
READ-based In-Memory Computing using Sentential Decision Diagrams . . . . .	818
ConvFIFO: A Crossbar Memory PIM Architecture for ConvNets Featuring First-In-First-Out Dataflow	824
<b>Session 8C: Firing Less for Evolution: Quantization &amp; Learning Spikes</b>	
MINT: Multiplier-less Integer Quantization for Energy Efficient Spiking Neural Networks . . . . .	830
TQ-TTFS: High-Accuracy and Energy-Efficient Spiking Neural Networks Using Temporal Quantization Time-to-First-Spike Neuron . . . . .	836



TEAS: Exploiting Spiking Activity for Temporal-wise Adaptive Spiking Neural Networks . . . . .	842
SOLSA: Neuromorphic Spatiotemporal Online Learning for Synaptic Adaptation . . . . .	848
<b>Session 8D: New Techniques for Photonics and Analog Circuit Design</b>	
Signed Convolution in Photonics with Phase-Change Materials using Mixed-Polarity Bitstreams . . . . .	854
An Efficient Branch-and-Bound Routing Algorithm for Optical NoCs . . . . .	860
Boosting Graph Spectral Sparsification via Parallel Sparse Approximate Inverse of Cholesky Factor . . . . .	866
Asynchronous Batch Constrained Multi-Objective Bayesian Optimization for Analog Circuit Sizing . . . . .	872
<b>Session 9A: Advancing AI Algorithms: Faster, Smarter, and More Efficient</b>	
Quantization-aware Optimization Approach for CNNs Inference on CPUs . . . . .	878
TSTC: Enabling Efficient Training via Structured Sparse Tensor Compilation . . . . .	884
An automated approach for improving the inference latency and energy efficiency of pretrained CNNs by removing irrelevant pixels with focused convolutions . . . . .	890
PIONEER: Highly Efficient and Accurate Hyperdimensional Computing using Learned Projection . . . . .	896
<b>Session 9B: Design Explorations for Neural Network Accelerators</b>	
Logic Design of Neural Networks for High-Throughput and Low-Power Applications . . . . .	902
Exact Scheduling to Minimize Off-Chip Data Movement for Deep Learning Accelerators . . . . .	908
Run-time Non-uniform Quantization for Dynamic Neural Networks in Wireless Communication . . . . .	915
PipeFuser: Building Flexible Pipeline Architecture for DNN Accelerators via Layer Fusion . . . . .	921
A Precision-Scalable RISC-V DNN Processor with On-Device Learning Capability at the Extreme Edge . . . . .	927
<b>Session 9C: High-Level Security Verification and Efficient Implementation</b>	
Microscope: Causality Inference Crossing the Hardware and Software Boundary from Hardware Per- spective . . . . .	933
d-GUARD: Thwarting Denial-of-Service Attacks via Hardware Monitoring of Information Flow using Language Semantics in Embedded Systems . . . . .	939
Security Coverage Metrics for Information Flow at the System Level . . . . .	945
Theoretical Patchability Quantification for IP-Level Hardware Patching Designs . . . . .	951
Multiplierless Design of High-Speed Very Large Constant Multiplications . . . . .	957
<b>Session 9D: Routing</b>	
V-GR: 3D Global Routing with Via Minimization and Multi-Strategy Rip-up and Rerouting . . . . .	963
A Fast and Robust Global Router with Capacity Reduction Techniques . . . . .	969
A High Performance Detailed Router Based on Integer Programming with Adaptive Route Guides . . . . .	975