

2025 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED 2025)

**Reykjavik, Iceland
6-8 August 2025**



**IEEE Catalog Number: CFP25LOW-POD
ISBN: 979-8-3315-2711-2**

**Copyright © 2025 by the Institute of Electrical and Electronics Engineers, Inc.
All Rights Reserved**

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

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

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

IEEE Catalog Number:	CFP25LOW-POD
ISBN (Print-On-Demand):	979-8-3315-2711-2
ISBN (Online):	979-8-3315-2710-5
ISSN:	1533-4678

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

Session 1A: Energy Efficient Computing and Storage

Chair: Frank K. Gurkaynak, ETH Zurich

[Partial-Sum Quantization Based on Pseudo-Quantization Noise for Variation-Tolerant Analog In-Memory Computing...1](#)

Nameun Kang, EunHyeok Park, Sangsu Park, Jongil Kim, Jaeyun Yi, Jae-Joon Kim

[An Analog Multiplier Utilizing an Unconventional Bit-Weighting Scheme with Application to Neural Network Quantization...8](#)

Mehdi Kamal, Massoud Pedram

[LoRASense: Learnable Low-Rank Acquisition in Sensors for Efficient Edge Machine Vision...15](#)

Yiwen Liang, Zhiqiang Yi, Tianrui Ma, Weidong Cao

[AdaGray: An Energy-Efficient Adaptive Gray-Code Strategy for QLC Flash-Memory Storage Systems...22](#)

Han-Yu Liao, Jen-Wei Hsieh, Yi-Shen Chen, Chang-Lin Tsai, Yuan-Hao Chang

Session 1B: Next-Gen EDA: Physical Design Optimization and AI-Driven Methodologies

Chair: Francesco Regazzoni, University of Amsterdam

[Simultaneous Optimization of Placement Legalization and Multi-bit Flip-flop Allocation in Physical Design Automation...28](#)

Seoyoung Bang, Taewhan Kim

[Timing-Driven Macro Placement with Connectivity-Aware Clustering...35](#)

Gangmin Jeon, Heechun Park

[GenSoC: A Multi-Agent-Assisted SoC Generation Methodology Leveraging Open-Source Hardware...42](#)

Peiran Yan, Qinzhe Zhi, Lifeng Liu, Tianyu Jia

[Development of a Physics-Informed Neural Network Model for Rapid Power Integrity Analysis in Die-Level and Die-Package Co-Design for 2.5-D Chiplet Solutions...49](#)

Xi Chen, Yuhao Ju, Jie Gu

Session 2A: Brains on a Budget – Smart, Small, and Power-Efficient

Chair: Hussam Amrouch, Technical University of Munich

[Exploration of Low-Power Flexible Stress Monitoring Classifiers for Conformal Wearables...56](#)

Florentia Afentaki, Sri Sai Rakesh Nakkilla, Konstantinos Balaskas, Paula Carolina Lozano Duarte, Shiyi Jiang, Georgios Zervakis, Farshad Firouzi, Krishnendu Chakrabarty, Mehdi Tahoori

[ASAP-FE: Energy-Efficient Feature Extraction Enabling Multi-Channel Keyword Spotting on Edge Processors...63](#)

Jongin Choi, Jina Park, Jae-Jin Lee, Massoud Pedram and Woojoo Lee

[Minimizing Redundant Checkpoint Triggers for Efficient Intermittent Systems...70](#)

Youngbin Kim, Yoojin Lim

[AridWalk: Efficient Graph Random Walks on a Resource-Limited Computational Storage Device...77](#)

Liang-Chi Chen, Chien-Chung Ho, Tei-Wei Kuo, Yuan-Hao Chang

Session 2B: Next-Gen Edge AI: Flash Memory, RISC-V, and Hybrid Architectures

Chair: Thierry Tambe, Stanford University

[E-Flash: Energy-Efficient DNN Mapping on NAND Flash Memory with State-Switching Algorithm...84](#)

Gisan Ji, Sanghun Shin, Jangho Baik, Wonbo Shim, Sungju Ryu

[Hybrid Systolic Array Accelerator with Optimized Dataflow for Edge Large Language Model Inference...91](#)

Chun-Ting Chen, HanGyeol Mun, Jian Meng, Mohamed Abdelfattah and Jae-sun Seo

[ID-VSA: Independent and Dynamic Vector Symbolic Architecture for Energy-Efficient Edge AI...98](#)

Mehran Shoushtari Moghadam, Abu Kaisar Mohammad Masum, Sercan Aygun, M. Hassan Najafi

[RIMIX: RISC-V Core with MIXed-Precision SIMD Instruction Extensions Supported by Oracle-Assisted Sub-Network Search for Efficient TinyML...105](#)

Jiyong Park, Dahoon Park, Yeeun Hong, Jaeha Kung

Session 3A: Accelerators for LLM and Edge AI Applications

Chair: M. Hassan Najafi, Case Western Reserve University

[Accelerating LLM Inference with Flexible N:M Sparsity via A Fully Digital Compute-in-Memory Accelerator...112](#)

Akshat Ramachandran, Souvik Kundu, Arnab Raha, Shamik Kundu, Deepak K. Mathaikutty, Tushar Krishna

[A Scalable External Memory Access and On-Chip Storage Architecture for Edge-AI Accelerators – Multi-Path Rolling Data Refresh and Layer-Wise Bank Allocation –...119](#)

Quan Cheng, Huizi Zhang, Qiufeng Li, Yuan Liang, Mingtao Zhang, Zhenzhe Chen, Ruilin Zhang, Jinjun Xiong, Mingqiang Huang, Longyang Lin, Masanori Hashimoto

FLASH-D: FlashAttention with Hidden Softmax Division...126

Kosmas Alexandridis, Vasileios Titopoulos, Giorgos Dimitrakopoulos

A Compact, Low Power Transprecision ALU for Smart Edge Devices...133

Ayushi Dube, Gian Singh, Sarma Vrudhula

Session 3B: Twisting AI and Hardware in 3D

Chair: Bokyung Kim, Rutgers University

MIX-3D: AI-based Architecture-Circuit Co-design Methodology for Mixed-Node, Mixed-Area 3D ICs...141

Min Gyu Park, Doyun Kim, Sung Kyu Lim

GAVINA: flexible aggressive undervolting for bit-serial mixed-precision DNN acceleration...148

Jordi Fornet, Pau Fontova-Musté, Adrian Gras, Omar Lahyani, Martí Caro, Jaume Abella, Francesc Moll Echeto, Josep Altet

An Open-Source HW-SW Co-Development Framework Enabling Efficient Multi-Accelerator Systems...155

Ryan Antonio, Joren Dumoulin, Xiaoling Yi, Josse Van Delm, Yunhao Deng, Guilherme Paim, Marian Verhelst

MemRaptor: Magnetoresistive Array as Matrix Vector Multiplication and Transcendental Function Operator for NLP Applications...162

Dong Eun Kim, Tanvi Sharma, Anushka Mukherjee, Mainakh Mukherjee and Kaushik Roy

Session 4A: Precision, Performance, and Efficiency: Modern Strategies in ML Systems

Chair: Donghwa Shin, Soongsil University

Faster Ternary and Binary Neural Network Inference on CPU by Reducing Popcount...169

Olivier Fischer, Shien Zhu, Gustavo Alonso

ECLIP: Energy-efficient and Practical Co-Location of ML Inference on Spatially Partitioned GPUs...176

Ryan Quach, Yidi Wang, Ali Jahanshahi, Daniel Wong, Hyoseung Kim

SmartMS: Efficient Hierarchical Database Search for Mass Spectrometry...183

Flavio Ponzina, Sumukh Pinge, Zheyu Li, Abhijay Deevi, Yilin Ge, Mingu Kang, Tajana Rosing

TruncQuant: Truncation-Ready Quantization for DNNs with Flexible Weight Bit Precision...190

Jinhee Kim, Seoyeon Yoon, Taeho Lee, Joo Chan Lee, Kang Eun Jeon, Jong Hwan Ko

Session 4B: Compute-in-Memory Macros

Chair: Jaehyun Park, University of Ulsan

CAM-CIM: A Hybrid Compute-in-Memory Using Content-Addressable Memory with Subword Split Mapping for Reduced ADC Resolution...197

Sangwoo Jung, Hojin Lee, Yejin Lee, Jiyong Park, Dahoon Park, Hyunseob Shin, Jong-Hyeok Yoon, Jaeha Kung

A 20.98TOPS/W Energy-Efficient Binary BERT Model on Group Vector Systolic CIM Accelerator...204

Dingbang Liu, Ziyi Guan, Qilong Chen, Jingyun Gu, Jiaqi Yang, Kai Li, Wei Mao, Ngai Wong, Changwen Chen, Hao Yu

Cost-efficient Processing-in-Memory Architecture with Training-free and Universal Error Compensation...211

Myeongji Yun, Jung Gyu Min, Sein Oh, Jiwoong Choi, Minkyu Je, Jang-Sik Lee, Youngjoo Lee

Design Techniques for Ultra-low Power Cryogenic CMOS for Quantum Computing Applications (Industry)...218

Sudipto Chakraborty, Pat Rosno, John Bulzacchelli, David Frank, Rajiv Joshi, Daniel Friedman

Session 5A: Transceivers and Stochastic Logic Design

Chair: Alexandre Levisse, EPFL

A Spectral-Efficient Low-Power NRZ/PAM-4 Dual-Mode Wireline Transmitter for Multidrop Interfaces...225

Donggeon Kim, Kiarash Gharibdoust, Armin Tajalli, Kyungtae Lee, Gain Kim

Energy-Efficient Single-Ended Capacitive PAM-4 Transceiver for Next-Generation HBM Interfaces...231

Jaeyoon Kim, Sangyoon Lee, Jaekwang Yun, Sanghyuk Seo, Kwangyeon Lee, Yong-Un Jeong, Suhwan Kim

Always-On Sensing in Energy-Harvested Systems via Stochastic Intermittent Computing...237

Sepehr Tabrizchi, Mehran Shoushtari Moghadam, Ali Shafiee Sarvestani, Sercan Aygun, M. Hassan Najafi, Arman Roohi

Design of a correlation-insensitive HFQ stochastic adder by local two-phase clocking...244

Yuki Matsumoto, Masamitsu Tanaka, Takatsugu Ono

Session 5B: Emerging Technologies for Secure and Efficient Computing

Chair: Tianyu Jia, Peking University

DPIMA: A DRAM-Based Processing-in-Memory Accelerator for Privacy-Preserving Machine Learning...250

Bokyung Kim

QuAKE: Speeding up Model Inference Using Quick and Approximate Kernels for Exponential Non-Linearities...257

Sai Kiran Narayanaswami, Gopalakrishnan Srinivasan, Ravindran Balaraman

[Meta-Heuristic Optimization of Karatsuba Multiplier Designed ECC Processor...264](#)

Pruthvi Parate, Daksh Sharma, Alwin Shaju, Vasanthi D R, Madhav Rao

[Unicorn-CIM: Uncovering the Vulnerability and Improving the Resilience of High-Precision Compute-in-Memory...271](#)

Qiufeng Li, Yiwen Liang, Weidong Cao

Special Session 1: Cold War: Don't Close Your Eyes on Temperature

Chair: Mircea Stan, University of Virginia

[Power Map Characterization and Modeling for Commercial CPU/GPUs Considering Temperature Dependence...429](#)

Jincong Lu, Sachin Sachdeva, Haotian Lu, Sheldon X.-D. Tan

[Thermal Challenges and Opportunities for Off-the-shelf 3D-stacked CPUs...436](#)

Jae Yoon Lee, Chae Young Sim, Seung Hun Choi, Sung Woo Chung

[Transistor-to-GDS Reliability Analysis in Sub-3nm: Impact of Self-Heating and Aging on Timing...443](#)

Swati Deshwal, Hadi Nour Eddine, Mahdi Benkhelifa, Albi Mema, Yogesh S. Chauhan, Hussam Amrouch

[Thermal Aware Design Methodologies for System On Chip Application Processor...449](#)

Youngsang Cho, Jun So Pak, Seungwook Yoon, Ilryong Kim

Special Session 2: Neuromorphic Edge Computing: Challenges, Opportunities, and Current Solutions

Chairs: Federico Corradi, Eindhoven University of Technology & Farhad Merchant, University of Groningen

[Dynamic Neuromorphic Processing for Energy-Efficient Cognitive Sensing...N/A](#)

Amir Zjajo

[Harnessing Sparsity for Low-Power Event-Driven Computing at the Edge...N/A](#)

Orlando Moreira

[Reliable and Low-Power Neuromorphic Computing...N/A](#)

Miloš Krstić

[Low-Power Neuromorphic Systems: Mixed-Signal Design and Computing-in-Memory for Edge AI...456](#)

Farhad Merchant, Federico Corradi

Tutorial 1: The Energy Cost of Privacy and Security

Chair: Francesco Regazzoni, University of Amsterdam

[Introduction to Security Primitives and Privacy Preserving Technologies...N/A](#)

Paolo Palmieri

[Energy Assessment of Security Primitives...463](#)

Ayşe Kivilcim Coskun

[Energy Efficient Design and Implementation of Security Primitives...N/A](#)

Francesco Regazzoni

Tutorial 2: Autonomy with Neuromorphic System

Chair: Saibal Mukhopadhyay, Georgia Institute of Technology

[Neuromorphic Circuits...464](#)

Amit Trivedi

[Neuromorphic and Hybrid Computing Models...N/A](#)

Priyadarshini Panda

[Algorithms for Neuromorphic Sensing Systems...N/A](#)

Saibal Mukhopadhyay

[Neuromorphic Approach to End-to-end Navigation...N/A](#)

Kaushik Roy

Poster Session

[On-chip Integrated Voltage Regulators: Frontside, Backside, or Off-Chip?...278](#)

Amaan Rahman, Seungmin Woo, Zheng Yang, Sung Kyu Lim

[J3DAI: A tiny DNN-Based Edge AI Accelerator for 3D-Stacked CMOS Image Sensor...285](#)

Benoit Tain, Raphael Millet, Romain Lemaire, Michal Szczepanski, Laurent Alacoque, Emmanuel Pluchart, Sylvain Choynet, Rohit Prasad, Jerome Chossat, Pascal Pierunek, Pascal Vivet, Sebastien Thuries

[MTA-Coded PAM-4 Receiver with Decision Feedback Power Saving Scheme and Partial DFE for Low-Power Memory Interfaces...292](#)

Jusung Lee, Younghwan Chang, Jaekwang Yun, Sanghyuk Seo, Yong-Un Jeong, Suhwan Kim

A High-Performance Dataflow-Based ORB Extractor Accelerator for SLAM...298

Rui Xue, Wenming Li, Haibin Wu, Cheng Guo, Yi Li, Xiaochun Ye, Dongrui Fan

Jack Unit: An Area- and Energy-Efficient Multiply-Accumulate (MAC) Unit Supporting Diverse Data Formats...305

Seock-Hwan Noh, Sungju Kim, Seohyun Kim, Daehoon Kim, Jaeha Kung, Yeseong Kim

Repurpose Accel-Sim for Next Generation NVIDIA Jetson GPU Architectural Design...312

Tianhao Huang, Lingyu Sun, Chao Li, Xiaofeng Hou, Yaqian Zhao, Jingwen Leng, Li Li, Minyi Guo

ML-Power: Machine Learning based Power Estimation for SoCs...319

Sujay Pandit, Sujit Dey, Anand Raghunathan

Diffusion-Enhanced Graph Transformer with Reinforcement Learning for Transferable Analog Circuit Optimizer...326

Ho-Jin Lee, Kyeong-Jun Lee, Jaehoon Lee, Kyu-Jin Choi, Geunyoung Choi, Youngchang Choi, Kyongsu Lee, Seokhyeong Kang, Jae-Yoon Sim

SHIFT ECC: A Value Converting HBM ECC Approach for Refresh Energy Efficient Integer Quantized DNN Inference...332

Jae Yoon Lee, Young Seo Lee, Young-Ho Gong, Seon Wook Kim, Sung Woo Chung

Optimizing Heterogeneous Compute-in-Memory with Hybrid Dataflow and In-Network Reduction for Vision Transformer...339

Zexin Fu, Yihang Zuo, Yuzhe Ma, Jiayi Huang

Revisiting Reconfigurable Acceleration of Vision Transformer with Patch Pruning...346

Hanning Chen, Yang Ni, Wenjun Huang, Hyunwoo Oh, Tamoghno Das, Fei Wen, Mohsen Imani

TEE-SFL: Time and Energy-efficient solution for addressing communication heterogeneity in Split Federated Learning Schemes...353

Ziyi Zhao, Qifeng Chen, Jingtao Li, Deliang Fan, Chaitali Chakrabarti

MEbots: Integrating a RISC-V Virtual Platform with a Robotic Simulator for Energy-aware Design...360

Giovanni Pollo, Mohamed Amine Hamdi, Matteo Rizzo, Lorenzo Ruotolo, Pietro Furbatto, Matteo Isoldi, Yukai Chen, Alessio Burrello, Enrico Macii, Massimo Poncino, Daniele Jahier Pagliari, Sara Vinco

Enhancing Low-Precision Deep Learning: A Posit8 Framework for Energy Efficient DNN Training...367

Dongyang Wu, Mehdi Kamal, Massoud Pedram

Towards Zero-Stall Matrix Multiplication on Energy-Efficient RISC-V Clusters for Machine Learning Acceleration...374

Luca Colagrande, Lorenzo Leone, Maximilian Coco, Andrei Deaconeasa, Luca Benini

Efficient Precision-Scalable Hardware for Microscaling (MX) Processing in Robotics Learning...381

Stef Cuyckens, Xiaoling Yi, Nitish Satya Murthy, Chao Fang, Marian Verhelst

SITRA: Exploiting Temporal Silence in Spiking Transformers for Fast & Energy-efficient Inference...388

Abhiroop Bhattacharjee, Abhishhek Moitra, Ruokai Yin, Priyadarshini Panda

DIRC-RAG: Accelerating Edge RAG with Robust High-Density and High-Loading-Bandwidth Digital In-ReRAM Computation...395

Kunming Shao, Zhipeng Liao, Jiangnan Yu, Liang Zhao, Qiwei Li, Xijie Huang, Jingyu He, Fengshi Tian, Yi Zou, Xiaomeng WANG, Tim Cheng, Chi Ying Tsui

upGEMV: A Bandwidth-Efficient and Scalable GEMV Accelerator for PIM Systems...402

Fan Yang, Shunchen Shi, Peiheng Zhang, Xueqi Li

Can Photonic Interconnects be used for High-Throughput Memory Access in FHE Accelerators?...409

Dewan Saiham, Mariam Rabadi, Di Wu, Sazadur Rahman

Sustainably Secure: ChaCha20 Encryption Based on In-Memory Compute...416

Samridhi Jain, Mohd Aamir, Anuj Grover

Design Contest Session

Chairs: Rajesh Kedia, IITH & Arnab Raha, Intel

ECO: Low Power Context-Aware Multimodal AI on NPUs...423

Arghadip Das, Yatharth Agarwal, Arnab Raha, Soumendu Ghosh, Vijay Raghunathan

Radar-PIM-Lite: Ultra-Low-Power PIM Processor for Real-Time UWB Radar Respiration Detection on UAVs...425

Kyeongwon Lee, Hyunseok Kwak, Kyeongpil Min, Chaebin Jung, Sangmin Jeon, Woojoo Lee, Jina Park, Massoud Pedram

A Low-Power Real-Time Hardware Accelerator for Edge Detection Using Stochastic Computing...427

Priyajit Ghosh, Rajarshi Mukherjee, Auro Anand Saha, Sutirtha Naha, Arghadip Das, Arnab Raha, Mrinal K Naskar