2023 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED 2023)

Vienna, Austria
7-8 August 2023

IEEE Catalog Number: CFP23LOW-POD
ISBN: 979-8-3503-1176-1

****** 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:	CFP23LOW-POD
ISBN (Print-On-Demand):	979-8-3503-1176-1
ISBN (Online):	979-8-3503-1175-4

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

Energy-Efficient Machine Learning Acceleration: from Technologies to Circuits and Systems 1
Chukwufumnanya Ogbogu, Madeleine Abernot, Corentin Delacour, Aida Todri-Sanial, Sudeep Pasricha, Partha Pratim Pande

A Comparative Study on Front-Side, Buried and Back-Side Power Rail Topologies in 3nm Technology Node 9
Sandra Maria Shaji, Lingjun Zhu, Junsik Yoon, Sung Kyu Lim

CoolDRAM: an Energy-Efficient and Robust DRAM 15
Nezam Rohbani, Mohammad Arman Soleimani, Hamid Sarbazi-Azad

Development of Tropical Algebraic Accelerator with Energy Efficient Time-Domain Computing for Combinatorial Optimization and Machine Learning 21
Qiankai Cao, Xi Chen, Jie Gu

IMBUE: In-Memory Boolean-to-CUrrent Inference ArchitecturE for Tsetlin Machines 27
Omar Ghazal, Simranjeet Singh, Tousif Rahman, Shengqi Yu, Yujin Zheng, Domenico Balsamo, Sachin Patkar, Farhad Merchant, Fei Xia, Alex Yakovlev, Rishad Shafik

IMAT: Energy-Efficient In-Memory Acceleration for Ternary Neural Networks with Sparse Dot Product 33
Shien Zhu, Shuo Huai, Guochu Xiong, Weichen Liu

TensorCV: Accelerating Inference-Adjacent Computation Using Tensor Processors 39
Dongho Ha, Won Woo Ro, Hung-Wei Tseng

Energy-Harvesting-Aware Adaptive Inference of Deep Neural Networks in Embedded Systems 45
Gwanjong Park, Osama Khan, Euiseong Seo

Precision-Aware Latency and Energy Balancing on Multi-Accelerator Platforms for DNN Inference 51
Matteo Risso, Alessio Burrello, Giuseppe Maria Sarda, Luca Benini, Enrico Macii, Massimo Poncino, Marian Verhelst, Daniele Jahier Pagliari

Florian: Developing a Low-Power RISC-V Multicore Processor with a Shared Lightweight FPU 57
Jina Park, Kyuseung Han, Eunjin Choi, Sukho Lee, Jae-Jin Lee, Woojoo Lee, Massoud Pedram

Energy Efficient Real-Time Scheduling on Heterogeneous Architectures with Self-Suspension 63
Wenwen Xu, Zheyu Zhang, Yuankai Xu, Jing Li, Yehan Ma, Yier Jin, Christopher D. Gill, Xuan Zhang, An Zou

CARMA: Context-Aware Runtime Reconfiguration for Energy-Efficient Sensor Fusion 69
Yifan Zhang, Arnav Vaibhav Malawade, Xiaofang Zhang, Yuhui Li, Donghwan Seong, Mohammad Abdullah Al Faruque, Sitao Huang

Uncertainty-Aware Online Learning for Dynamic Power Management in Large Manycore Systems 75
Gaurav Narang, Raid Ayoub, Michael Kishinevsky, Janardhan Rao Doppa, Partha Pratim Pande

A Multicore GNN Training Accelerator 81
Sudipta Mondal, Ramprasath S., Ziqing Zeng, Kishor Kunal, Sachin S. Sapatnekar

Joint Optimization of Cache Management and Graph Reordering for GCN Acceleration 87
Kyeong-Jun Lee, Byungjun Kim, Han-Gyeol Mun, Seunghyun Moon, Jae-Yoon Sim

ITA: An Energy-Efficient Attention and Softmax Accelerator for Quantized Transformers 93
Gamze Islamoglu, Moritz Scherer, Gianna Paulin, Tim Fischer, Victor J. B. Jung, Angelo Garofalo, Luca Benini

Energy-Efficient RISC-V-Based Vector Processor for Cache-Aware Structurally-Pruned Transformers 99
Jung Gyu Min, Dongyun Kam, Younghoon Byun, Gunho Park, Youngjoo Lee

Machine Learning Driven Synthesis of Clock Gating 105
Doyeon Won, Soomin Kim, Taewhan Kim

Automatic Generation of Structured Macros Using Standard Cells – Application to CIM 111
Christian Lanius, Jie Lou, Johnson Loh, Tobias Gemmeke

Multi-Objective Optimization for Floating Point Mix-Precision Tuning 117
Zeqing Li, Yongwei Wu, Youhui Zhang

REFROM: Responsive, Energy-Efficient Frame Rendering for Mobile Devices 123
Tsung-Yen Hsu, Yi-Shen Chen, Yun-Chih Chen, Yuan-Hao Chang, Tei-Wei Kuo

DCIM-3DRec: A 3D Reconstruction Accelerator with Digital Computing-in-Memory and Octree-Based Scheduler 129
Yiqi Jing, Yiyang Sun, Xiao Wang, Wentao Zhao, Meng Wu, Fengyun Yan, Yufei Ma, Le Ye, Tianyu Jia

Processing-In-Memory Using Optically-Addressed Phase Change Memory 135
Guowei Yang, Cansu Demirkiran, Zeynep Ece Kizilates, Carlos A. Ríos Ocampo, Ayse K. Coskun, Ajay Joshi

LAXOR: A Bit-Accurate BNN Accelerator with Latch-XOR Logic for Local Computing 141
Dongrui Li, Tomomasa Yamasaki, Aarthy Mani, Anh Tuan Do, Niangjun Chen, Bo Wang

AR-PIM: An Adaptive-Range Processing-in-Memory Architecture 147
Teyuh Chou, Fernando Garcia-Redondo, Paul Whatmough, Zhengya Zhang

Low Power Logic Obfuscation Through System Level Clock Gating 153
Daniel Xing, Yuntao Liu, Ankur Srivastava

FPGA-Patch: Mitigating Remote Side-Channel Attacks on FPGAs Using Dynamic Patch Generation 159
Mahya Morid Ahmadi, Lilas Alrahis, Ozgur Sinanoglu, Muhammad Shafique

Enabling DVFS Side-Channel Attacks for Neural Network Fingerprinting in Edge Inference Services 165
Erich Malan, Valentino Peluso, Andrea Calimera, Enrico Macii

Hardware Trojans in fdSOI 171
Christian Lanius, Florian Freye, Shutao Zhang, Tobias Gemmeke

Bridging the Gap Between Spiking Neural Networks & LSTMs for Latency & Energy Efficiency 177
Gourav Datta, Haoqin Deng, Robert Aviles, Zeyu Liu, Peter A. Beerel

Partial-Sum Quantization for Near ADC-Less Compute-In-Memory Accelerators 183
Utkarsh Saxena, Kaushik Roy

Efficient Multi-Objective Optimization for PVT Variation-Aware Circuit Sizing Using Surrogate Models and Smart Corner Sampling 189
Octavian Pascu, Catalin Visan, Georgian Nicolae, Mihai Boldeanu, Horia Cucu, Cristian Diaconu, Andi Buzo, Georg Pelz

Model-Driven Dataset Generation for Data-Driven Battery SOH Models 195
Khaled Sidahmed Sidahmed Alamin, Francesco Daghero, Giovanni Pollo, Daniele Jahier Pagliari, Yukai Chen, Enrico Macii, Massimo Poncino, Sara Vinco

Ocellus: Highly Parallel Convolution-In-Pixel Scheme Realizing Power-Delay-Efficient Edge Intelligence 201
Sepehr Tabrizchi, Shaahin Angizi, Arman Roohi

Sky-NN: Enabling Efficient Neural Network Data Processing with Skyrmion Racetrack Memory 207
Yong-Cheng Liaw, Shuo-Han Chen, Yuan-Hao Chang, Yu-Pei Liang

RF2P: A Lightweight RISC Processor Optimized for Rapid Migration from IEEE-754 to Posit 213
Hyun Woo Oh, Seongmo An, Won Sik Jeong, Seung Eun Lee

Scaled Population Division for Approximate Computing 219
Kunal Bharathi, Sunil P. Khatri, Jiang Hu

Cryogenic CMOS as an Enabler for Low Power Dynamic Logic 225
Rakshith Saligram, Suman Datta, Arijit Raychowdhury

Quantifying the Overheads of Modular Multiplication 231
Deepraj Soni, Mohammed Nabeel, Negar Neda, Ramesh Karri, Michail Maniatakos, Brandon Reagen

Multi-Source Transfer Learning for Design Technology Co-Optimization 237
Jakang Lee, Jaeseung Lee, Seonghyeon Park, Seokhyeong Kang

Enabling Highly-Efficient DNA Sequence Mapping Via ReRAM-based TCAM 243
Yu-Shao Lai, Shuo-Han Chen, Yuan-Hao Chang

A Self-Powered Predictive Maintenance System Based on Piezoelectric Energy Harvesting and TinyML 249
Zijie Chen, Yiming Gao, Junrui Liang

Temperature-Aware Memory Mapping and Active Cooling of Neural Processing Units 255
Vahidreza Moghaddas, Hammam Kattan, Tim Bücher, Mikail Yayla, Jian-Jia Chen, Hussam Amrouch

WeNet: Configurable Neural Network with Dynamic Weight-Enabling for Efficient Inference 261
Jingxiao Ma, Sherief Reda

Energy-Efficient Missing Data Recovery in Wearable Devices: a Novel Search-Based Approach 267
Dina Hussein, Taha Belkhouja, Ganapati Bhat, Janardhan Rao Doppa

RecPIM: A PIM-Enabled DRAM-RRAM Hybrid Memory System for Recommendation Models 273
Heewoo Kim, Haojie Ye, Trevor Mudge, Ronald Dreslinski, Nishil Talati

Weight-Aware Activation Mapping for Energy-Efficient Convolution on PIM Arrays 279
Kang Eun Jeon, Johnny Rhe, Hyeonsu Bang, Jong Hwan Ko

Teleport: A High-Performance ShiftNet Hardware Accelerator with Fused Layer Computation 285
Hyunmin Kim, Sungju Ryu

Energy-Efficient ReRAM-Based ML Training Via Mixed Pruning and Reconfigurable ADC 291
Chukwufumnanya Ogbogu, Mohapatra Soumen, Biresh Kumar Joardar, Janardhan Rao Doppa, Deuk Heo, Krishnendu Chakrabarty, Partha Pratim Pande

Digital Implementation of On-Chip Hebbian Learning for Oscillatory Neural Network 297
Edgar Luhulima, Madeleine Abernot, Federico Corradi, Aida Todri-Sanial

PAIRS: Pruning-AIded Row-Skipping for SDK-Based Convolutional Weight Mapping in Processing-In-Memory Architectures 303
Johnny Rhe, Kang Eun Jeon, Jong Hwan Ko

A Fully-Integrated Energy-Scalable Transformer Accelerator Supporting Adaptive Model Configuration and Word Elimination for Language Understanding on Edge Devices 309
Zexi Ji, Hanrui Wang, Miaorong Wang, Win-San Khwa, Meng-Fan Chang, Song Han, Anantha P. Chandrakasan

Learning from Output Transitions: A Chosen Challenge Strategy for ML Attacks on PUFs 315
Chia-Chih Lin, Ming-Syan Chen

Efficient Machine Learning on Encrypted Data Using Hyperdimensional Computing 321
Yujin Nam, Minxuan Zhou, Saransh Gupta, Gabrielle De Micheli, Rosario Cammarota, Chris Wilkerson, Daniele Micciancio, Tajana Rosing

Author Index