

2025 35th International Conference on Field-Programmable Logic and Applications (FPL 2025)

**Leiden, Netherlands
1-5 September 2025**

IEEE Catalog Number:
ISBN:

CFP25623-POD
979-8-3315-9154-0



Copyright © 2025, IEEE

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:	CFP25623-POD
ISBN (Print-On-Demand):	979-8-3315-9154-0
ISBN (Online):	979-8-3315-9153-3
ISSN:	1946-147X

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

Virtualization and Dynamic Reconfiguration of Custom Instruction Accelerators (CIA) in RISC-V Embedded Systems	1
<i>Bea Healy, Brandon Freiburger, Jonas Kuentler, King Lok Chung, Emil Cozac, Meinhard Kissich Gennadiy Knis, Ron Sass, Dirk Koch, Jan Gray, Guy Lemieux</i>	
GEF: A GNN-Based Evaluation Framework for FPGA Routing Architecture	10
<i>Yuanqi Wang, Yunfei Dai, Jiangnan Li, Kaixiang Zhu, Huizhen Kuang, Hao Zhou, Eric Ren, Xifan Tang Weijun Qin, Tao Li, Lingli Wang</i>	
FPGA-Based MPSoCs for High-Performance Sensor Fusion: Accelerating Covariance Intersection	19
<i>Hazem M. Sharf, Mohamed Hassan</i>	
ATAPP: Architecture and Technology Aware Power Predictor for Unseen FPGAS	29
<i>Zhigang Wei, Aman Arora, Emily Shriver, Lizy K. John</i>	
Programmable Congestion Generator for Evaluating FPGA Interconnect Robustness	39
<i>Atreyee Saha, Sandesh Goyal, Aashish Tripathi</i>	
Routino: Accelerating FPGA Routing Through Efficient Memory Representation	48
<i>Davide Nicolini, Corrado De Sio, Eleonora Vacca, Luca Sterpone</i>	
Open-Source FPGA Routing Runtime Prediction for Improved Productivity Via Smart Route Termination	57
<i>Andrew David Gunter, Steven J E Wilton</i>	
Energy-Efficient DNNs on FPGAs for Edge-Cloud Computer Vision	66
<i>Qaisar Farooq, Idilio Drago</i>	
Towards Accelerated Healthcare Federated System Through Heterogeneous Accelerators	68
<i>Giuseppe Sorrentino, Davide Conficconi</i>	
Using Data to Reduce Uncertainty in FPGA Routing	70
<i>Andrew David Gunter, Steven J E Wilton</i>	
Reducing FPGA Placement Runtime by Clustering of Netlist Blocks	72
<i>Markus Rein, Dirk Stroobandt</i>	
Interconnection-Aware Resynthesis for Improving FPGA Physical Design	74
<i>Xiaoke Wang, Dirk Stroobandt</i>	
Multi-FPGA Programming Using OpenMP	76
<i>Pedro Henrique Rosso, Guido Araujo</i>	
FPGAs with FABulous – Framework and Chips	77
<i>Dirk Koch, Myrtle Shah, King Lok Chung, Jonas Kuentler, Marcel Jung, Jakob Ternes, Asma Mohsin Gennadiy Knis</i>	
NPX: Automating Neuromorphic Processor Design from Spike-Based Learning to FPGA Prototyping	78
<i>Kyuseung Han, Hyeonuk Jang, Sukho Lee, Sung-Eun Kim, Kyudong Hwang, Jae-Jin Lee</i>	

Live Demonstration: Continuous Processing of Event-Data with Graph Convolutional Neural Networks Implemented for SoC FPGA	79
<i>Piotr Wzorek, Kamil Jeziorek, Marcin Kowalczyk, Krzysztof Blachut, Tomasz Kryjak, Marek Gorgon</i>	
Versatile Place and Route with Continuous Routing Runtime Prediction and Smart Route Termination	80
<i>Andrew David Gunter, Steven J E Wilton</i>	
Compile in Seconds and Run on an FPGA with DynaRapid	81
<i>Andrea Guerrieri, Isaac John Wetenkamp, Chris Lavin, Eddie Hung, Lana Josipović, Paolo Ienne</i>	
ReconFormer: A Multi-Level Run-Time Reconfigurable System-on-Chip for Accelerating Transformers	82
<i>Je Yang, Gabriele Tombesi, Joseph Zuckerman, Luca P. Carloni</i>	
FLAIC: A Novel FPGA Logic Architecture via Fine-Grained Cut Topology Analysis	90
<i>Xianfeng Cao, Huizhen Kuang, Yuanqi Wang, Lingli Wang</i>	
Double Duty: FPGA Architecture to Enable Concurrent LUT and Adder Chain Usage	99
<i>Junius Pun, Xilai Dai, Grace Zgheib, Mahesh A. Iyer, Andrew Boutros, Vaughn Betz Mohamed S. Abdelfattah</i>	
DEFA: Design Space Exploration for FPGA Overlay Accelerators Through Frequency Prediction and Bayesian Optimization	109
<i>Qilong Zhu, Yunfei Dai, Shiyan Bi, Huizhen Kuang, Dylan Wang, Wenbo Yin, Lingli Wang</i>	
FINN-GL: Generalized Mixed-Precision Extensions for FPGA-Accelerated LSTMS	118
<i>Shashwat Khandelwal, Jakoba Petri-Koenig, Thomas B. Preußner, Michaela Blott, Shanker Shreejith</i>	
Multiplexer Optimizations for Virtex FPGAs	127
<i>Nicolai Fiege, Martin Hardieck, Peter Zipf</i>	
Identifying Sat Resilient Blocks Through LUT Switching Analysis for Breaking Compound Logic Locking Schemes	132
<i>Binbin Wang, Xinmu Wang, Xinyu Zhang, Shibo Tang, Huisi Zhou, Wei Hu</i>	
URAM-Based Asynchronous FIFO Design for Improved Throughput and FPGA RAM Usage	140
<i>Martim Rosado, Pedro Tomás, Nuno Roma, André David</i>	
From Errors to Solutions: LLM-Powered Command Scripting for FPGA Cad Tools	145
<i>Mohamed A. Elgammal, Vaughn Betz</i>	
TRPlaceFPGA-MP: A Two-Stage Reinforcement Learning Framework for Fast FPGA Macro Placer	155
<i>Qin Luo, Xinshi Zang, Evangeline F.Y. Young, Martin D.F. Wong</i>	
FAME: FPGA Acceleration of Secure Matrix Multiplication with Homomorphic Encryption	163
<i>Zhihan Xu, Rajgopal Kannan, Viktor K. Prasanna</i>	
F ² Opt: Novel Fine-Tuning and Folding Algorithms for FPGA-Based DNN Accelerators	173
<i>Muhammad Shakeel Akram, Bogaraju Sharatchandra Varma, Vincent Meyers, Mehdi Tahoori Dewar Finlay</i>	
AffiNiTy: A Multi-Scalar Multiplication Accelerator with a Novel Batched Inversion Architecture	183
<i>Tong Wu, Niall Emmart, Oliver Diessel</i>	

EViL: An Efficient Vision-LSTM Accelerator Based on FPGA	192
<i>Zexuan Deng, Han Jiao, Wenjin Huang, Yihua Huang</i>	
SparseDPD: A Sparse Neural Network-Based Digital Predistortion FPGA Accelerator for RF Power Amplifier Linearization	197
<i>Manno Versluis, Yizhuo Wu, Chang Gao</i>	
Routing Struggle: A Metric to Quantify Routability	202
<i>Sadegh Yazdanshenas, Jeffrey Chromczak</i>	
Aspo: Constraint-Aware Bayesian Optimization for Fpga-Based Soft Processors	211
<i>Haoran Wu, Ce Guo, Wayne Luk, Robert Mullins</i>	
Improving Boolean Satisfiability-Based Modulo Scheduling	220
<i>Nicolai Fiege, Peter Zipf</i>	
Maximum FPGA: A 32K-Point 32-Parallel Floating Point FFT	229
<i>Martin Langhammer, Bogdan Pasca</i>	
Cocotb-Pynq: Co-Simulating Python+RTL Applications Targeting Pynq Platforms with Cocotb	238
<i>Gavin Lusby, Nachiket Kapre</i>	
EQViTA: an End-To-End Quantized Vision Transformer Accelerator Implemented on Resource-Constrained FPGAs	243
<i>Jiacheng Cao, Jiaqi Guo, Wei Xiong, Huanlin Luo, Jian Wang, Jinmei Lai</i>	
FPGA Stereo Visual Slam with Efficient Stereo Feature Matching and Key-Frame Generation	252
<i>Miyuru Thathsara, Damith Anhettigama, Siew-Kei Lam</i>	
A High-Performance and Resource-Efficient FPGA-Based Multi-Object Tracking System Using Event Cameras	261
<i>Wei Xiong, Jianfan Zhang, Xingzhe Zhu, Jiacheng Cao, Jian Wang, Jinmei Lai</i>	
Four-Input Lookup Table (LUT4) and Architectural Enhancements Enable Power Efficient Mid-Range FPGAs	270
<i>Satwant Singh, Michael Schneider, Ziad Aboud, Jonathan Peterson, Senani Gunaratna, Ting Yew Cindy Lee, Rick Crotty</i>	
Refining Datapath for Microscaling ViTs	279
<i>Can Xiao, Jianyi Cheng, Yiren Zhao</i>	
Towards Instruction-Controlled In-Pipeline GEMM Acceleration in a Dual-Issue RISC-V Core for Edge Applications	289
<i>Ahmad Othman, Darmen Ilyas, Ahmed Kamaleldin, Diana Göhringer</i>	
Design Space Exploration of Fast RISC-V Processors for Scalable Kilo-Core FPGA Systems	298
<i>Riadh Ben Abdelhamid, Vladislav Valek, Kevin Klein, Dirk Koch</i>	
DMA Calypste: Open-Source Ultra-Low Latency DMA Engine for FPGAs	307
<i>Vladislav Válek, Martin Špinler, Jakub Cabal, Tomáš Martínek</i>	
Hybrid Weightless Neural Networks for Efficient Edge Inference	312
<i>Mugdha Jadhao, Alan T. L. Bacellar, Shashank Nag, Igor D. S. Miranda, Felipe M. G. França, Lizy K. John</i>	

Accelerating Transposed Convolutions on FPGA-Based Edge Devices	321
<i>Jude Haris, José Cano</i>	
AMD Versal Implementations of FAM and SSCA Estimators	330
<i>Carol Jingyi Li, Ruilin Wu, Philip H.W. Leong</i>	
GAMA: High-Performance GEMM Acceleration on AMD Versal ML-Optimized AI Engines	339
<i>Kaustubh Manohar Mhatre, Endri Taka, Aman Arora</i>	
NeuGEMM: A Reordering-Free Unified GEMM-Conv2D Accelerator for Lightweight Neuromorphic Processors	348
<i>Hyeonuk Jang, Sukho Lee, Jae-Jin Lee, Kyuseung Han</i>	
Accelerating K-Means: A Vectorized Approach for AI Engines & Neural Processing Units	353
<i>Eleonora Cabai, Giuseppe Sorrentino, Marco Domenico Santambrogio, Davide Conficconi</i>	
Maximizing Resource Utilization for Stencil Computing	358
<i>Pedro Henrique Rosso, Guido Araujo</i>	
Performance Impact on Reducing Energy Consumption Applying Adaptive Stream-Based Entropy Coding on FPGA	360
<i>Taku Nishikawa, Koichi Marumo, Shinichi Yamagiwa</i>	
Analytical Buffer Sizing for Neural Network Inference Applications on FPGAs	361
<i>Lukas Stasytis, Felix Jentzsch, Zsolt István</i>	
Connection Tracking at 400 Gbps	362
<i>David Vodák, Oliver Gurka, Jiří Matoušek, Daniel Kondys</i>	