

2024 34th International Conference on Field-Programmable Logic and Applications (FPL 2024)

**Torino, Italy
2-6 September 2024**



**IEEE Catalog Number: CFP24623-POD
ISBN: 979-8-3315-3008-2**

**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:	CFP24623-POD
ISBN (Print-On-Demand):	979-8-3315-3008-2
ISBN (Online):	979-8-3315-3007-5
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

CURRAN ASSOCIATES INC.
proceedings
.com

2024 34th International Conference on Field- Programmable Logic and Applications (FPL) **FPL 2024**

Table of Contents

Message from the General and Program Chairs	xii
Organizing Committee	xiii
Program Committee	xiv
Steering Committee	xvii
Reviewers	xviii
Keynotes	xxi
Tutorials	xxiv
Sponsors	xxx

Session 1: Applications

Exploring the Versal AI Engines for Signal Processing in Radio Astronomy	1
<i>Victor van Wijhe (Netherlands Institute for Radio Astronomy, The Netherlands; University of Twente, The Netherlands), Vincent Sprave (Otto-von-Guericke University Magdeburg, Germany), Daniele Passaretti (Otto-von-Guericke University Magdeburg, Germany), Nikolaos Alachiotis (University of Twente, The Netherlands), Gerrit Grutzeck (Max Planck Institute for Radio Astronomy, Germany), Thilo Pionteck (Otto-von-Guericke University Magdeburg, Germany), and Steven van der Vlugt (Netherlands Institute for Radio Astronomy, The Netherlands)</i>	
JSON-CooP: A JSON Decompression/Parsing Co-Design for FPGAs	11
<i>Tobias Hahn (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany), Stefan Wildermann (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany), and Jürgen Teich (Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Germany)</i>	
KIT: Kernel Isotropic Transformation of Bilateral Filters for Image Denoising on FPGA	19
<i>Fanny Spagnolo (University of Calabria, Italy), Pasquale Corsonello (University of Calabria, Italy), Fabio Frustaci (University of Calabria, Italy), and Stefania Perri (University of Calabria, Italy)</i>	

Session 2: Placement & Routing

DynaRapid: Fast-Tracking from C to Routed Circuits	24
<i>Andrea Guerrieri (EPFL, Switzerland), Srijeet Guha (EPFL, Switzerland), Chris Lavin (AMD Research and Advanced Development), Eddie Hung (AMD Research and Advanced Development), Lana Josipović (ETH Zurich, Switzerland), and Paolo Ienne (EPFL, Switzerland)</i>	
The Road Less Traveled: Congestion-Aware NoC Placement and Packet Routing for FPGAs	33
<i>Soheil Gholami Shahrouz (University of Toronto, Canada) and Vaughn Betz (University of Toronto, Canada)</i>	
Better Together: Combining Analytical and Annealing Methods for FPGA Placement	43
<i>Rachel Selina Rajarathnam (The University of Texas at Austin, USA), Kate Thurmer (University of Toronto, Canada), Vaughn Betz (University of Toronto, Canada), Mahesh A. Iyer (Intel Corporation, USA), and David Z. Pan (The University of Texas at Austin, USA)</i>	
A High-Performance Routing Engine for Large-Scale FPGAs	53
<i>Timothy Martin (University of Guelph, Canada), Dani Maarouf (University of Guelph, Canada), Gary Grewal (University of Guelph, Canada), and Shawki Areibi (University of Guelph, Canada)</i>	

Session 3: High-Bandwidth & Virtual Memory

SERI: High-Throughput Streaming Acceleration of Electron Repulsion Integral Computation in Quantum Chemistry using HBM-Based FPGAs	60
<i>Philip Stachura (Simon Fraser University, Canada), Guanyu Li (Simon Fraser University, Canada), Xin Wu (Paderborn Center for Parallel Computing; Paderborn University, Germany), Christian Plessl (Paderborn Center for Parallel Computing; Paderborn University, Germany), and Zhenman Fang (Simon Fraser University, Canada)</i>	
H2PIPE: High Throughput CNN Inference on FPGAs with High-Bandwidth Memory	69
<i>Mario Doumet (University of Toronto, Canada), Marius Stan (Microsoft, USA), Mathew Hall (Microsoft, USA), and Vaughn Betz (University of Toronto, Canada; Vector Institute, Canada)</i>	
FlexiMem: Flexible Shared Virtual Memory for PCIe-Attached FPGAs	78
<i>Canberk Sönmez (École Polytechnique Fédérale de Lausanne, Switzerland), Mohamed Shahawy (École Polytechnique Fédérale de Lausanne, Switzerland), Cemalettin Cem Belentepe (École Polytechnique Fédérale de Lausanne, Switzerland), and Paolo Ienne (École Polytechnique Fédérale de Lausanne, Switzerland)</i>	
SoGraph: A State-Aware Architecture for Out-of-Memory Graph Processing on HBM-Equipped FPGAs	87
<i>Qianyu Cheng (University of Science and Technology of China, China), Zhendong Zheng (University of Science and Technology of China, China), Tianhao Jiang (University of Science and Technology of China, China), Cheng Tang (University of Science and Technology of China, China), Teng Wang (University of Science and Technology of China, China), Lei Gong (University of Science and Technology of China, China), Chao Wang (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)</i>	

Leveraging HBM2 for Accelerating k-mer Counting with oneAPI on FPGAs	92
<i>Owen P. Lucas (ECE, University of Pittsburgh NSF SHREC Center, USA)</i>	
<i>and Alan D. George (ECE, University of Pittsburgh NSF SHREC Center, USA)</i>	

Session 4: High-Level Synthesis & Simulation

StencilStream: A SYCL-Based Stencil Simulation Framework Targeting FPGAs	100
<i>Jan-Oliver Opdenhövel (Paderborn University), Christoph Alt (Paderborn University; Friedrich Alexander Universität, Erlangen), Christian Plessl (Paderborn University), and Tobias Kenter (Paderborn University)</i>	
Efficient Design Space Exploration for Dynamic & Speculative High-Level Synthesis	109
<i>Dylan Leothaud (Univ Rennes, Inria, CNRS, IRISA), Jean-Michel Gorius (Univ Rennes, Inria, CNRS, IRISA), Simon Rokicki (Univ Rennes, Inria, CNRS, IRISA), and Steven Derrien (Univ Rennes, Inria, CNRS, IRISA)</i>	
Fast Switching Activity Estimation for HLS-Produced Dataflow Circuits	118
<i>Jiantao Liu (ETH Zurich, Switzerland), Maksymilian Graczyk (ETH Zurich, Switzerland), Andrea Guerrieri (EPFL, Switzerland & HES-SO, Switzerland), and Lana Josipović (ETH Zurich, Switzerland)</i>	
FlexWalker: An Efficient Multi-Objective Design Space Exploration Framework for HLS Design....	126
<i>Zheyuan Zou (University of Science and Technology of China, China), Cheng Tang (School of Software Engineering), Lei Gong (University of Science and Technology of China, China), Chao Wang (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)</i>	
Chimera: A co-Simulation Framework Combining with gem5 and FPGA Platform for Efficient Verification	133
<i>Chao Fu (Fudan University, China), Zengshi Wang (Fudan University, China), and Jun Han (Fudan University, China)</i>	

Session 5: Machine Learning 1

NeuraLUT: Hiding Neural Network Density in Boolean Synthesizable Functions	140
<i>Marta Andronic (Imperial College London, UK) and George A. Constantinides (Imperial College London, UK)</i>	
PolyLUT-Add: FPGA-Based LUT Inference with Wide Inputs	149
<i>Binglei Lou (The University of Sydney, Australia), Richard Rademacher (The University of Sydney, Australia), David Boland (The University of Sydney, Australia), and Philip H.W. Leong (The University of Sydney, Australia)</i>	
Kratos: An FPGA Benchmark for Unrolled DNNs with Fine-Grained Sparsity and Mixed Precision	156
<i>Xilai Dai, Yuzong Chen, and Mohamed S. Abdelfattah</i>	

Session 6: Cryptography & Security

UniGuard: A Unified Hardware-Oriented Threat Detector for FPGA-Based AI Accelerators	164
<i>Xiaobei Yan (Nanyang Technological University, Singapore), Han Qiu (Tsinghua University, China), and Tianwei Zhang (Nanyang Technological University, Singapore; National Integrated Centre for Evaluation (NiCE), Singapore)</i>	
A Better Kyber Butterfly for FPGAs	171
<i>Jonas Bertels (COSIC, Belgium), Quinten Norga (COSIC, Belgium), and Ingrid Verbauwhede (COSIC, Belgium)</i>	
Techniques for Exploring Fine-Grained LUT and Routing Aging on a 28nm FPGA	178
<i>Hayden Cook (Brigham Young University, USA) and Jeffrey Goeders (Brigham Young University, USA)</i>	

Session 7: Architectures

A Software-Programmable Neural Processing Unit for Graph Neural Network Inference on FPGAs.....	187
<i>Taikun Zhang (University of Toronto, Canada; Programmable Solutions Group, Intel Corporation), Andrew Boutros (University of Toronto, Canada), Sergey Gribok (Programmable Solutions Group, Intel Corporation), Kwadwo Boateng (Programmable Solutions Group, Intel Corporation), and Vaughn Betz (University of Toronto, Canada)</i>	
Revealing Untapped DSP Optimization Potentials for FPGA-Based Systolic Matrix Engines	197
<i>Jindong Li (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence; University of Chinese Academy of Sciences), Tenglong Li (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence; University of Chinese Academy of Sciences), Guobin Shen (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence; School of Future Technology), Dongcheng Zhao (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence), Qian Zhang (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence; University of Chinese Academy of Sciences), and Yi Zeng (Brain-inspired Cognitive Intelligence Lab, Institute of Automation, Chinese Academy of Sciences; Center for Long-term Artificial Intelligence; School of Future Technology; University of Chinese Academy of Sciences; Key Laboratory of Brain Cognition and Brain-inspired Intelligence Technology, Chinese Academy of Sciences)</i>	
SA4: A Comprehensive Analysis and Optimization of Systolic Array Architecture for 4-bit Convolutions	204
<i>Geng Yang (Xidian University), Jie Lei (University of Technology Sydney), Zhenman Fang (Simon Fraser University), Jiaqing Zhang (Xidian University), Junrong Zhang (Xidian University), Weiying Xie (Xidian University), and Yunsong Li (Xidian University)</i>	

CFEACT: A CGRA-Based Framework Enabling Agile CNN and Transformer Accelerator Design ..	213
<i>Yiqing Mao (Fudan University, China), Xuchen Gao (Fudan University, China), Jiahang Lou (Fudan University, China), Yunhui Qiu (Fudan University, China), Wenbo Yin (Fudan University, China), Wai-Shing Luk (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
IMAGine: An In-Memory Accelerated GEMV Engine Overlay	220
<i>Arafat Kabir (University of Arkansas), Tendayi Kamucheka (University of Arkansas), Nathaniel Fredricks (University of Arkansas), Joel Mandebi (Advanced Micro Devices, Inc.), Jason Bakos (University of South Carolina), Miaoqing Huang (University of Arkansas), and David Andrews (University of Arkansas)</i>	

Session 8: Machine Learning 2

AMA: An Analytical Approach to Maximizing the Efficiency of Deep Learning on Versal AI Engine	227
<i>Xiaodong Deng (Xi'an Jiaotong University), Shijie Wang (Xi'an Jiaotong University), Tianyi Gao (Xi'an Jiaotong University), Jing Liu (Xi'an Jiaotong University), Longjun Liu (Xi'an Jiaotong University), and Nanning Zheng (Xi'an Jiaotong University)</i>	
A Heterogeneous Acceleration System for Attention-Based Multi-Agent Reinforcement Learning..	236
<i>Samuel Wiggins (University of Southern California), Yuan Meng (University of Southern California), Mahesh A. Iyer (Intel Corporation), and Viktor Prasanna (University of Southern California)</i>	
Fitop-Trans: Maximizing Transformer Pipeline Efficiency through Fixed-Length Token Pruning on FPGA	243
<i>Kejia Shi (Fudan University, China), Manting Zhang (Fudan University, China), Keqing Zhao (Fudan University, China), Xiaoxing Wu (Fudan University, China), Yang Liu (Fudan University, China), Jun Yu (Fudan University, China), and Kun Wang (Fudan University, China)</i>	
An Open-Source and Extensible Framework for Fast Prototyping and Benchmarking of Spiking Neural Network Hardware	250
<i>Shadi Matinizadeh (Drexel University, USA) and Anup Das (Drexel University, USA)</i>	
HASS: Hardware-Aware Sparsity Search for Dataflow DNN Accelerator	257
<i>Zhewen Yu (Imperial College London, UK), Sudarshan Sreeram (Imperial College London, UK), Krish Agrawal (Imperial College London, UK), Junyi Wu (Imperial College London, UK), Alexander Montgomerie-Corcoran (Imperial College London, UK), Cheng Zhang (Imperial College London, UK), Jianyi Cheng (University of Cambridge, UK), Christos-Savvas Bouganis (Imperial College London, UK), and Yiren Zhao (Imperial College London, UK)</i>	

Session 9: Edge & Low-Power Computing

SDA: Low-Bit Stable Diffusion Acceleration on Edge FPGAs	264
<i>Geng Yang (Xidian University; Simon Fraser University), Yanyue Xie (Northeastern University), Zhong Jia Xue (Simon Fraser University), Sung-En Chang (Northeastern University), Yanyu Li (Northeastern University), Peiyan Dong (Northeastern University), Jie Lei (University of Technology Sydney), Weiyang Xie (Xidian University), Yanzhi Wang (Northeastern University), Xue Lin (Northeastern University), and Zhenman Fang (Simon Fraser University)</i>	
E^3HDC: Energy Efficient Encoding for Hyper-Dimensional Computing on Edge Devices	274
<i>Mahboobe Sadeghipour Roodsari (Institute of Computer Engineering, Karlsruhe Institute of Technology, Germany), Jonas Krautter (Institute of Computer Engineering, Karlsruhe Institute of Technology, Germany), Vincent Meyers (Institute of Computer Engineering, Karlsruhe Institute of Technology, Germany), and Mehdi Tahoori (Institute of Computer Engineering, Karlsruhe Institute of Technology, Germany)</i>	
Energy-Aware Synchronization of Hardware Tasks in Virtualized Embedded Systems	281
<i>Cornelia Wulf (Technische Universität Dresden, Germany), Gökhan Akgün (Technische Universität Dresden, Germany), Mehdi Safarpour (University of Oulu, Finland), Anastacia Grishchenko (Technische Universität Dresden, Germany), and Diana Göhringer (Technische Universität Dresden, Germany)</i>	

Session 10: Arithmetic

FPGA Modular Multipliers using Hybrid Reduction Techniques	288
<i>Sergey Gribok (Intel Corporation, United States), Martin Langhammer (Intel Corporation, United Kingdom), and Bogdan Pasca (Intel Corporation, France)</i>	
Shedding the Bits: Pushing the Boundaries of Quantization with Minifloats on FPGAs	297
<i>Shivam Aggarwal (National University of Singapore, Singapore; AMD Research, Ireland), Hans Jakob Damsgaard (Electrical Engineering Unit, Tampere University, Finland; AMD Research, Ireland), Alessandro Pappalardo (National University of Singapore, Singapore), Giuseppe Franco (Electrical Engineering Unit, Tampere University, Finland), Thomas B. Preußner (AMD Research, Ireland), Michaela Blott (AMD Research, Ireland), and Tulika Mitra (National University of Singapore, Singapore)</i>	
Exploring FPGA Designs for MX and Beyond	304
<i>Ebby Samson (Imperial College London, UK), Naveen Mellempudi (AMD, USA), Wayne Luk (Imperial College London, UK), and George A. Constantinides (Imperial College London, UK)</i>	
Fast and Practical Strassen's Matrix Multiplication using FPGAs	311
<i>Afzal Ahmad (The Hong Kong University of Science and Technology, Hong Kong), Linfeng Du (The Hong Kong University of Science and Technology, Hong Kong), and Wei Zhang (The Hong Kong University of Science and Technology, Hong Kong)</i>	

Session 11: Accelerators

FORC: A High-Throughput Streaming FPGA Accelerator for Optimized Row Columnar File Decoders in Big Data Engines	318
<i>Abdul Wadood (Simon Fraser University, Canada), Alec Lu (Simon Fraser University, Canada), Ken Zhang (Huawei Canada, Canada), and Zhenman Fang (Simon Fraser University, Canada)</i>	
BitBlender: Scalable Bloom Filter Acceleration on FPGAs with Dynamic Scheduling	325
<i>Kenneth Liu (HiAccel Lab, Simon Fraser University, Canada), Alec Lu (HiAccel Lab, Simon Fraser University, Canada), and Zhenman Fang (HiAccel Lab, Simon Fraser University, Canada)</i>	
LORA: A Latency-Oriented Recurrent Architecture for GPT Model on Multi-FPGA Platform with Communication Optimization	332
<i>ZhenDong Zheng (University of Science and Technology of China, China), Qianyu Cheng (University of Science and Technology of China, China), Teng Wang (University of Science and Technology of China, China), Lei Gong (University of Science and Technology of China, China), Xianglan Chen (University of Science and Technology of China, China), Cheng Tang (University of Science and Technology of China, China), Chao Wang (University of Science and Technology of China, China), and Xuehai Zhou (University of Science and Technology of China, China)</i>	
DTrans: A Dataflow-Transformation FPGA Accelerator with Nonlinear-Operators Fusion Aiming for the Generative Model	339
<i>Xuanzheng Wang (Tsinghua University), Shuo Miao (Tsinghua University), Peng Qu (Tsinghua University), and Youhui Zhang (Tsinghua University)</i>	
CFSA: An Efficient CPU-FPGA Synergies Accelerator for Neural Radiation Field Rendering	346
<i>Shangrong Li (Xidian University, China), Kai Liu (Xidian University, China), Wei Liu (Smart Earth Key Laboratory, China), and Zibo Guo (Xidian University, China)</i>	
Author Index	353