

2023 International Conference on Field Programmable Technology (ICFPT 2023)

**Yokohama, Japan
12-14 December 2023**



**IEEE Catalog Number: CFP23528-POD
ISBN: 979-8-3503-5912-1**

**Copyright © 2023 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:	CFP23528-POD
ISBN (Print-On-Demand):	979-8-3503-5912-1
ISBN (Online):	979-8-3503-5911-4
ISSN:	2837-0430

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

2023 International Conference on Field Programmable Technology (ICFPT) **ICFPT 2023**

Table of Contents

Message from General Chair and Program Co-Chairs	xii
Organizing Committee	xiv
Program Committee	xv
Steering Committee	xviii
Artifact Evaluation Committee	xix
Sponsors and Supporters	xx

Journal Track Papers

On the Malicious Potential of Xilinx' Internal Configuration Access Port (ICAP)	1
Covert-channels in FPGA-enabled SmartSSDs	2
HyBNN: Quantifying and Optimizing Hardware Efficiency of Binary Neural Networks	3
Across Time and Space: Senju's Approach for Scaling Iterative Stencil Loop Accelerators on Single and Multiple FPGAs	5
AEKA: FPGA Implementation of Area-Efficient Karatsuba Accelerator for Ring-Binary-LWE-based Lightweight PQC	6

Full Paper

A Deep-Learning Data-Driven Approach for Reducing FPGA Routing Runtimes	7
<i>T. Martin (University of Guelph, Canada), Q. Li (University of Guelph, Canada), C. Barnes (University of Guelph, Canada), G. Grewal (University of Guelph, Canada), and S. Areibi (University of Guelph, Canada)</i>	
A Tenant Side Compilation Solution for Cloud FPGA Deployment	16
<i>Maximillian Panoff (University of Florida, USA), Muhammed Kawser Ahmed (University of Florida, USA), Hanqiu Wang (University of Florida, USA), Shuo Wang (University of Florida, USA), and Christophe Bobda (University of Florida, USA)</i>	
GRAFT: GNN-Based Adaptive Framework for Efficient CGRA Mapping	26
<i>Jiangnan Li (Fudan University, China), Chang Cai (Fudan University, China), Yaya Zhao (Renmin University of China, China), Yazhou Yan (Fudan University, China), Wenbo Yin (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	

BSTMSM: A High-Performance FPGA-Based Multi-Scalar Multiplication Hardware Accelerator	35
<i>Baoze Zhao (Sun Yat-sen University, China), Wenjin Huang (Sun Yat-sen University, China), Tianrui Li (Sun Yat-sen University, China), and Yihua Huang (Sun Yat-sen University, China; Southern Marine Science and Engineering Guangdong Laboratory, China)</i>	
FPGA-Accelerated Quantum Transport Measurements	44
<i>Timo Haarman (University of Twente, The Netherlands), Antonio Sousa de Almeida (University of Twente, The Netherlands), Amber Heskes (University of Twente, The Netherlands), Floris Zwanenburg (University of Twente, The Netherlands), and Nikolaos Alachiotis (University of Twente, The Netherlands)</i>	
An Efficient Dataflow for Convolutional Generative Models	53
<i>Zhengzheng Ma (Peking University) and Guojie Luo (Peking University)</i>	
PolyLUT: Learning Piecewise Polynomials for Ultra-Low Latency FPGA LUT-Based Inference	60
<i>Marta Andronic (Imperial College London, UK) and George A. Constantinides (Imperial College London, UK)</i>	
M4BRAM: Mixed-Precision Matrix-Matrix Multiplication in FPGA Block RAMs	69
<i>Yuzong Chen (Cornell University), Jordan Dotzel (Cornell University), and Mohamed S. Abdelfattah (Cornell University)</i>	
VIB: A Versatile Interconnection Block for FPGA Routing Architecture	79
<i>Kaichuang Shi (Fudan University, China), Hao Zhou (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
AUGER: A Multi-Objective Design Space Exploration Framework for CGRAs	88
<i>Jingyuan Li (Fudan University, China), Yihan Hu (Fudan University, China), Yuan Dai (Fudan University, China), Huizhen Kuang (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
MaxEVA: Maximizing the Efficiency of Matrix Multiplication on Versal AI Engine	96
<i>Endri Taka (The University of Texas at Austin, USA), Aman Arora (Arizona State University, USA, The University of Texas at Austin, USA), Kai-Chiang Wu (National Yang Ming Chiao Tung University, Taiwan, The University of Texas at Austin, USA), and Diana Marculescu (The University of Texas at Austin)</i>	
HGBO-DSE: Hierarchical GNN and Bayesian Optimization Based HLS Design Space Exploration .	106
<i>Huizhen Kuang (Fudan University, China), Xianfeng Cao (Fudan University, China), Jingyuan Li (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
A High-Frequency Load-Store Queue with Speculative Allocations for High-Level Synthesis	115
<i>Robert Szafarczyk (University of Glasgow, UK), Syed Waqar Nabi (University of Glasgow, UK), and Wim Vanderbauwhede (University of Glasgow, UK)</i>	
Efficient FPGA-Based Accelerator for Post-Processing in Object Detection	125
<i>Zibo Guo (Xidian University, China), Kai Liu (Xidian University, China), Wei Liu (State Key Laboratory of Geo-information Engineering, China), and Shangrong Li (Xidian University, China)</i>	

Extending Data Flow Architectures for Convolutional Neural Networks to Multiple FPGAs	132
<i>Mohamed Ibrahim (University of Toronto, Canada; Intel Corporation), Zhipeng Zhao (Carnegie Mellon University, United States), Mathew Hall (Microsoft, United States), and Vaughn Betz (University of Toronto, Canada; Vector Institute for AI)</i>	
Assuring Netlist-to-Bitstream Equivalence Using Physical Netlist Generation and Structural Comparison	142
<i>Reilly McKendrick (Brigham Young University, USA), Keenan Faulkner (Brigham Young University, USA), and Jeffrey Goeders (Brigham Young University, USA)</i>	
Respect the Difference: Reinforcement Learning for Heterogeneous FPGA Placement	152
<i>Fatemehsadat Mahmoudi (University of Toronto, Canada), Mohamed A. Elgammal (University of Toronto, Canada), Soheil Gholami Shahrouz (University of Toronto, Canada), Kevin E. Murray (Cerebras Systems, Canada), and Vaughn Betz (University of Toronto, Canada)</i>	
SSiMD: Supporting Six Signed Multiplications in a DSP Block for Low-Precision CNN on FPGAs..	161
<i>Qi Liu (Zhejiang University, China), Mo Sun (Zhejiang University, China), Jie Sun (Zhejiang University, China), Liqiang Lu (Zhejiang University, China), Jieru Zhao (Shanghai Jiao Tong University), and Zeke Wang (Zhejiang University, China)</i>	
LUTNet-RC: Look-Up Tables Networks for Reservoir Computing on an FPGA	170
<i>Kanta Yoshioka (Kyushu Institute of Technology, Japan), Yuichiro Tanaka (Kyushu Institute of Technology, Japan), and Hakaru Tamukoh (Kyushu Institute of Technology, Japan)</i>	
SATAY: A Streaming Architecture Toolflow for Accelerating YOLO Models on FPGA Devices	179
<i>Alexander Montgomerie-Corcoran (Imperial College London, UK), Petros Toupas (Imperial College London, UK), Zhewen Yu (Imperial College London, UK), and Christos-Savvas Bouganis (Imperial College London, UK)</i>	
Mercury: An Automated Remote Side-Channel Attack to Nvidia Deep Learning Accelerator	188
<i>Xiaobei Yan (Nanyang Technological University, Singapore), Xiaoxuan Lou (Nanyang Technological University, Singapore), Guowen Xu (City University of Hong Kong, China), Han Qiu (Tsinghua University, China), Shangwei Guo (Chongqing University, China), Chip Hong Chang (Nanyang Technological University, Singapore), and Tianwei Zhang (Nanyang Technological University, Singapore)</i>	
Into the Third Dimension: Architecture Exploration Tools for 3D Reconfigurable Acceleration Devices	198
<i>Andrew Boutros (University of Toronto, Canada), Fatemehsadat Mahmoudi (University of Toronto, Canada), Amin Mohaghegh (University of Toronto, Canada), Stephen More (University of Toronto, Canada), and Vaughn Betz (University of Toronto, Canada)</i>	
Integrated Multi-Ported Memory Distribution for Temporal-Multiplexing Workloads on FPGAs ...	209
<i>Chia-Chen Yen (National Taiwan University, Taiwan), Mi-Yen Yeh (Academia Sinica, Taiwan), and Ming-Syan Chen (National Taiwan University, Taiwan)</i>	

OD-REM: On-Demand Regular Expression Matching on FPGAs for Efficient Deep Packet Inspection	217
<i>Weihai Xu (Southeast University, China; Purple Mountain Laboratories, China), Zheng Zhou (Southeast University, China; Purple Mountain Laboratories, China), Jin Zhang (Purple Mountain Laboratories, China), Yiming Jiang (National Digital Switching Engineering; Technological Research Center, China), and Peng Yi (National Digital Switching Engineering; Technological Research Center, China)</i>	

Short Paper

Asymmetry in Butterfly Fat Tree FPGA NoC	227
<i>Dongjoon Park (University of Pennsylvania, USA), Zhijing Yao (University of Pennsylvania, USA), Yuanlong Xiao (University of Pennsylvania, USA), and André DeHon (University of Pennsylvania, USA)</i>	
Accelerated Real-Time Classification of Evolving Data Streams Using Adaptive Random Forests	232
<i>Frank Ridder (University of Twente, The Netherlands), Kuan-Hsun Chen (University of Twente, The Netherlands), and Nikolaos Alachiotis (University of Twente, The Netherlands)</i>	
SqueezeBlock: A Transparent Weight Compression Scheme for Deep Neural Networks	238
<i>Mo Song (The University of Hong Kong, Hong Kong), Jiajun Wu (The University of Hong Kong, Hong Kong), Yuhao Ding (The University of Hong Kong, Hong Kong), and Hayden Kwok-Hay So (The University of Hong Kong, Hong Kong)</i>	
Efficiently Removing Sparsity for High-Throughput Stream Processing	244
<i>Philippos Papaphilippou (Trinity College Dublin, Ireland; Imperial College London, UK), Zhiqiang Que (Imperial College London, UK), and Wayne Luk (Imperial College London, UK)</i>	

Ph.D. Forum

Explore the Feedback Interconnects in Intra-Cluster Routing for FPGAs	250
<i>Kaichuang Shi (Fudan University, China), Hao Zhou (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
An FPGA-GPU Heterogeneous System and Implementation for On-Board Remote Sensing Data Processing	254
<i>Tingting Qiao (Beijing Institute of Technology, China), Yu Xie (Beijing Institute of Technology, China), He Chen (Beijing Institute of Technology, China), and Yizhuang Xie (Beijing Institute of Technology, China)</i>	
An Extremely Pipelined FPGA-Based Accelerator of All Adder Neural Networks for On-Board Remote Sensing Scene Classification	258
<i>Ning Zhang (Beijing Institute of Technology, China), Shuo Ni (Beijing Institute of Technology, China), Tingting Qiao (Beijing Institute of Technology, China), Wenchao Liu (Beijing Institute of Technology, China), and He Chen (Beijing Institute of Technology, China)</i>	

Performance Modeling and Scalability Analysis of Stream Computing in ESSPER FPGA Clusters .. 262
Ryota Miyagi (The University of Tokyo, Japan), Ryota Yasudo (Kyoto University, Japan), Kentaro Sano (RIKEN, Japan), and Hideki Takase (The University of Tokyo, Japan)

Demo Night

Traffic Flow Optimization Using a Chaotic Boltzmann Machine Annealer on an FPGA 266
Kanta Yoshioka (Kyushu Institute of Technology, Japan), Yuichiro Tanaka (Kyushu Institute of Technology, Japan), and Hakaru Tamukoh (Kyushu Institute of Technology, Japan)

Kyokko: a Virtual Channel Capable Aurora 64B/66B Compatible Serial Communication Controller 268
Akinobu Tomori (Kumamoto University, Japan) and Yasunori Osana (Kumamoto University, Japan)

Offloading Image Recognition Processing for Care Robots to FPGA on Multi-Access Edge Computing 270
Hayato Mori (Shibaura Institute of Technology, Japan), Eisuke Okazaki (Tokai University, Japan), Gai Nagahashi (Tokai University, Japan), Mikiko Sato (Tokai University, Japan), Takeshi Ohkawa (Kumamoto University, Japan), and Midori Sugaya (Shibaura Institute of Technology, Japan)

A State Vector Quantum Simulator Working on FPGAs with Extensible SATA Storage 272
Kaijie Wei (Keio University, Japan), Ryohei Niwase (University of Tsukuba, Japan), Hideharu Amano (Keio University, Japan), Yoshiki Yamaguchi (University of Tsukuba, Japan), and Takefumi Miyoshi (WasaLabo, LLC., Japan)

Stochastic Implementation of Simulated Quantum Annealing on PYNQ 274
Taiga Kubota (Tohoku University, Japan), Duckgyu Shin (Tohoku University, Japan), Naoya Onizawa (Tohoku University, Japan), and Takahiro Hanyu (Tohoku University, Japan)

Poster

An FPGA-Based Mix-Grained Sparse Training Accelerator 276
Yingchang Mao (Tianjin University, China) and Qiang Liu (Tianjin University, China)

DSLUT: An Asymmetric LUT and its Automatic Design Flow Based on Practical Functions 278
Moucheng Yang (State Key Laboratory of ASIC and System, Fudan University, Shanghai, China), Kaixiang Zhu (State Key Laboratory of ASIC and System, Fudan University, Shanghai, China), Lingli Wang (State Key Laboratory of ASIC and System, Fudan University, Shanghai, China), and Xuegong Zhou (Institute of Big Data, Fudan University, Shanghai, China)

F-TFM: Accelerating Total Focusing Method for Ultrasonic Array Imaging on FPGA 280
Bizhao Shi (Peking University, China), Jieran Zhang (Peking University, China), and Guojie Luo (Peking University, China)

FPGA Resource-Aware Structured Pruning for Real-Time Neural Networks	282
<i>Benjamin Ramhorst (Imperial College London, United Kingdom), Vladimir Lončar (Massachusetts Institute of Technology, United States of America), and George A. Constantinides (Imperial College London, United Kingdom)</i>	
UAV Swarm Planning Accelerator on FPGA with low Latency and Fixed-Point L-BFGS Quasi-Newton Solver	284
<i>Suquan Zhang (Tsinghua University, China), Jincheng Yu (Tsinghua University, China), Yuanfan Xu (Tsinghua University, China), and Wang Yu (Tsinghua University, China)</i>	
FPGA Framework Improvements for HPC Applications	286
<i>Antonio Filgueras (Universitat Politècnica de Catalunya; Barcelona Supercomputing Center), Miquel Vidal (Barcelona Supercomputing Center), Daniel Jiménez-González (Universitat Politècnica de Catalunya; Barcelona Supercomputing Center), Carlos Álvarez (Universitat Politècnica de Catalunya; Barcelona Supercomputing Center), and Xavier Martorell (Universitat Politècnica de Catalunya; Barcelona Supercomputing Center)</i>	
Enormous-Scale Quantum State Vector Calculation with FPGA-Accelerated SATA Storages	288
<i>Ryohei Niwase (University of Tsukuba, Japan), Hikaru Harasawa (University of Tsukuba, Japan), Yoshiki Yamaguchi (University of Tsukuba, Japan), Kaijie Wei (Keio University, Japan), Hideharu Amano (Keio University, Japan), and Takefumi Miyoshi (WasaLabo, LLC., Japan)</i>	
Introducing the NAIL Accelerator Interface Layer for Low Latency FPGA Offload	290
<i>Edward Grindley (The Alan Turing Institute, United Kingdom), Thurstan Gray (The Alan Turing Institute, United Kingdom), James Wilkinson (The Alan Turing Institute, United Kingdom), Chris Vaux (The Alan Turing Institute, United Kingdom), Adam Ardron (The Alan Turing Institute, United Kingdom), Jack Deeley (The Alan Turing Institute, United Kingdom), Alexander Elliott (The Alan Turing Institute, United Kingdom), Nidhin Thandassery Sumithran (University of Warwick, United Kingdom), and Suhaib A. Fahmy (University of Warwick, United Kingdom)</i>	
Towards Asynchronously Triggered Spiking Neural Network on FPGA for Event-Based Vision	292
<i>Zhenyu Wu (The University of Hong Kong, Hong Kong), Mo Song (The University of Hong Kong, Hong Kong), and Hayden Kwok-Hay So (The University of Hong Kong, Hong Kong)</i>	
AMG: Automated Efficient Approximate Multiplier Generator for FPGAs via Bayesian Optimization	294
<i>Zhen Li (Fudan University, China), Hao Zhou (Fudan University, China), Lingli Wang (Fudan University, China), and Xuegong Zhou (Fudan University, China)</i>	
E ² -ACE: An Energy-Efficient Reconfigurable Crypto-Accelerator with Agile End-to-End Toolchain	296
<i>Yuhang Cao (Fudan University, China), Yunhui Qiu (Fudan University, China), XuChen Gao (Fudan University, China), Qilong Zhu (Fudan University, China), Wenbo Yin (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	

A Dynamic Partial Reconfigurable CGRA Framework for Multi-Kernel Applications	298
<i>Qilong Zhu (Fudan University, China), Yuhang Cao (Fudan University, China), Yunhui Qiu (Fudan University, China), Xuchen Gao (Fudan University, China), Wenbo Yin (Fudan University, China), and Lingli Wang (Fudan University, China)</i>	
DiffLo: A Graph-Based Method for Functional Discrepancy Localization in High-Level Synthesis	300
<i>Liangji Chen (Hong Kong University of Science and Technology, China), Tingyuan Liang (Hong Kong University of Science and Technology, China), Wei Zhang (Hong Kong University of Science and Technology, China), and Sharad Sinha (Indian Institute of Technology Goa, India)</i>	
OpenTitan Based Multi-Level Security in FPGA System-on-Chips	302
<i>Sujan Kumar Saha (University of Florida, Gainesville), Abigail N. Butka (University of Florida, Gainesville), Muhammed Kawser Ahmed (University of Florida, Gainesville), and Christophe Bobda (University of Florida, Gainesville)</i>	
Author Index	305