2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM 2020)

Fayetteville, Arkansas, USA
3 – 6 May 2020

IEEE

IEEE Catalog Number: CFP20054-POD
ISBN: 978-1-7281-5804-4

Copyright © 2020 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:	CFP20054-POD
ISBN (Print-On-Demand):	978-1-7281-5804-4
ISBN (Online):	978-1-7281-5803-7
ISSN:	2576-2613

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

2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)

FCCM 2020

Table of Contents

Message from the FCCM 2020 General and Program Chairs xiii
Organizing Committee xv
Program Committee xvi
Additional Reviewers xviii
Sponsors xix

Session 1: Machine Learning

High-Throughput Convolutional Neural Network on an FPGA by Customized JPEG Compression 1
Hiroki Nakahara (Tokyo Institute of Technology), Zhiqiang Que (Imperial College London), and Wayne Luk (Imperial College London)

Optimizing Reconfigurable Recurrent Neural Networks 10
Zhiqiang Que (Imperial College London), Hiroki Nakahara (Tokyo Institute of Technology), Eriko Nurvitadhi (Intel Corporation), Hongxiang Fan (Imperial College London), Chenglong Zeng (Corerain Technologies Ltd.), Jiuxi Meng (Imperial College London), Xinyu Niu (Corerain Technologies Ltd.), and Wayne Luk (Imperial College London)

Accelerating Proximal Policy Optimization on CPU-FPGA Heterogeneous Platforms 19
Yuan Meng (University of Southern California), Sanmukh Kuppannagari (University of Southern California), and Viktor Prasanna (University of Southern California)

Evaluating Low-Memory GEMMs for Convolutional Neural Network Inference on FPGAs 28
Wentai Zhang (Peking University), Ming Jiang (Peking University), and Guojie Luo (Peking University)

CNN-Based Feature-Point Extraction for Real-Time Visual SLAM on Embedded FPGA 33
Zhilin Xu (Tsinghua University), Jincheng Yu (Tsinghua University), Chao Yu (Tsinghua University), Hao Shen (Meituan-Dianping Group), Yu Wang (Tsinghua University), and Huazhong Yang (Tsinghua University)

Session 2: Networks and Security

Corundum: An Open-Source 100-Gbps NIC 38
Alex Forencich (University of California, San Diego), Alex C. Snoeren (University of California, San Diego), George Porter (University of California, San Diego), and George Papen (University of California, San Diego)

FFShark: A 100G FPGA Implementation of BPF Filtering for Wireshark 47
Juan Camilo Vega (University of Toronto), Marco Antonio Merlini (University of Toronto), and Paul Chow (University of Toronto)

Hardware Architecture of a Number Theoretic Transform for a Bootstrappable RNS-Based Homomorphic Encryption Scheme 56
Sunwoong Kim (University of Washington), Keewoo Lee (Seoul National University), Wonhee Cho (Seoul National University), Yujin Nam (Seoul National University), Jung Hee Cheon (Seoul National University), and Rob A. Rutenbar (University of Pittsburgh)

Power-Hammering through Glitch Amplification – Attacks and Mitigation 65
Kaspar Matas (The University of Manchester), Tuan Minh La (The University of Manchester), Khoa Dang Pham (The University of Manchester), and Dirk Koch (The University of Manchester)

Exploring the Impact of Switch Arity on Butterfly Fat Tree FPGA NoCs 70
Ian Lang (University of Waterloo), Ziqiang Huang (University of Waterloo), and Nachiket Kapre (University of Waterloo)

Session 3: Arithmetic

Comparison of Arithmetic Number Formats for Inference in Sum-Product Networks on FPGAs 75
Lukas Sommer (TU Darmstadt), Lukas Weber (TU Darmstadt), Martin Kumm (Fulda University of Applied Sciences), and Andreas Koch (TU Darmstadt)

High Density 8-Bit Multiplier Systolic Arrays for FPGA 84
Martin Langhammer (Intel Corporation), Sergey Gribok (Intel Corporation), and Gregg Baeckler (Intel Corporation)

Low-Cost Approximate Constant Coefficient Hybrid Binary-Unary Multiplier for DSP Applications 93
S. Rasoul Faraji (University of Minnesota, Twin Cities), Pierre Abillama (University of Minnesota, Twin Cities), and Kia Bazargan (University of Minnesota, Twin Cities)

Session 4: Virtualization, HBM, and Soft Processors

Enabling Efficient and Flexible FPGA Virtualization for Deep Learning in the Cloud 102
Shulin Zeng (Tsinghua University), Guohao Dai (Tsinghua University), Hanbo Sun (Tsinghua University), Kai Zhong (Tsinghua University), Guangjun Ge (Tsinghua University), Kaiyuan Guo (Tsinghua University), Yu Wang (Tsinghua University), and Huazhong Yang (Tsinghua University)

Shuhai: Benchmarking High Bandwidth Memory on FPGAs .111
Zeke Wang (Zhejiang University), Hongjing Huang (Zhejiang University), Jie Zhang (Zhejiang University), and Gustavo Alonso (ETH Zurich)

Exploring Writeback Designs for Efficiently Leveraging Parallel-Execution Units in FPGA-Based Soft-Processors .120
Eric Matthews (Simon Fraser University), Yuhui Gao (Simon Fraser University), and Lesley Shannon (Simon Fraser University)

Session 5: Applications

Safely Preventing Unbounded Delays During Bus Transactions in FPGA-Based SoC .129
Francesco Restuccia (Scuola Superiore Sant'Anna), Alessandro Biondi (Scuola Superiore Sant'Anna), Mauro Marinoni (Scuola Superiore Sant'Anna), and Giorgio Buttazzo (Scuola Superiore Sant'Anna)

Grapefruit: An Open-Source, Full-Stack, and Customizable Automata Processing on FPGAs .138
Reza Rahimi (University of Virginia), Elaheh Sadredini (University of Virginia), Mircea Stan (University of Virginia), and Kevin Skadron (University of Virginia)

FP-AMG: FPGA-Based Acceleration Framework for Algebraic Multigrid Solvers .148
Pouya Haghi (Boston University), Tong Geng (Boston University), Anqi Guo (Boston University), Tianqi Wang (University of Science and Technology of China), and Martin Herbordt (Boston University)

Algorithm-Hardware Co-Design for BQSR Acceleration in Genome Analysis ToolKit .157
Michael Lo (University of California, Los Angeles), Zhenman Fang (Simon Fraser University), Jie Wang (University of California, Los Angeles), Peipei Zhou (University of California, Los Angeles), Mau-Chung Frank Chang (University of California, Los Angeles), and Jason Cong (University of California, Los Angeles)

A Turbo Maximum-a-Posteriori Equalizer for Faster-Than-Nyquist Applications .167
Mohamed Omran Matar (University of British Columbia), Mrinmoy Jana (University of British Columbia), Jeebak Mitra (Huawei Canada), Lutz Lampe (University of British Columbia), and Mieszko Lis (University of British Columbia)

FPGA-Accelerated Automatic Alignment for Three-Dimensional Tomography .172
Shuang Wen (Peking University) and Guojie Luo (Peking University)

Session 6: HLS and Tooling

Artisan: A Meta-Programming Approach for Codifying Optimisation Strategies .177
Jessica Vandebon (Imperial College London), Jose G. F. Coutinho (Imperial College London), Wayne Luk (Imperial College London), Eriko Nurvitadhi (Intel Corporation), and Tim Todman (Imperial College London)

Hierarchical Modelling of Generators in Design-Space Exploration .186
Charles Lo (University of Toronto) and Paul Chow (University of Toronto)

Investigating Performance Losses in High-Level Synthesis for Stencil Computations .195
Wesson Altoyan (Stanford University) and Juan J. Alonso (Stanford University)

Poster Session 1: Arithmetic and Security

Proposing a Fast and Scalable Systolic Array for Matrix Multiplication .204
Bahar Asgari (Georgia Institute of Technology), Ramyad Hadidi (Georgia Institute of Technology), and Hyesoon Kim (Georgia Institute of Technology)

An Automated Tool for Design Space Exploration of Matrix Vector Multiplication (MVM) Kernels Using OpenCL Based Implementation on FPGAs .205
Jannatun Naher (North Carolina A & T State University), Clay Gloster (North Carolina A & T State University), Christopher C. Doss (North Carolina A & T State University), and Shrikanth S. Jadhav (North Carolina A & T State University)

Fast Arithmetic Hardware Library for RLWE-Based Homomorphic Encryption .206
Rashmi Agrawal (Boston University), Lake Bu (The Charles Stark Draper Laboratory), and Michel A. Kinsy (Boston University)

Primitive Instantiation for Speed-Area Efficient Architecture Design of Cellular Automata Based Mageto Logic on FPGA with Built-In Testability .207
Ayan Palchaudhuri (Indian Institute of Technology Kharagpur) and Anindya Sundar Dhar (Indian Institute of Technology Kharagpur)

TBOX-Based Mask Scrambling Against SCA .208
João Carlos Resende (Universidade de Lisboa), Ricardo J. R. Maçãs (Universidade de Lisboa), and Ricardo Chaves (Universidade de Lisboa)

FPGA Implementation of Post-Quantum DME Cryptosystem .209
José L. Imaña (Complutense University) and Ignacio Luengo (Complutense University)

A Dynamic Frequency Scaling Framework Against Reliability and Security Issues in Multi-tenant FPGA .210
Yukui Luo (University of Illinois at Chicago) and Xiaolin Xu (University of Illinois at Chicago)

Poster Session 2: Datacenter and Infrastructure

SHIP: Storage for Hybrid Interconnected Processors .211
Juan Camilo Vega (University of Toronto), Qianfeng (Clark) Shen (University of Toronto), and Paul Chow (University of Toronto)

RISC-V Barrel Processor for Accelerator Control .212
MohammadHossein AskariHemmat (Ecole Polytechnique Montreal), Olexa Bilaniuk (University of Montreal), Sean Wagner (IBM Canada), Yvon Savaria (Ecole Polytechnique Montreal), and Jean-Pierre David (Ecole Polytechnique Montreal)

Update Latency Optimization of Packet Classification for SDN Switch on FPGA .213
Chenglong Li (National University of Defense Technology), Tao Li (National University of Defense Technology), Junnan Li (National University of Defense Technology), Zilin Shi (National University of Defense Technology), and Baosheng Wang (National University of Defense Technology)

Accommodating Multi-tenant FPGAs in the Cloud .214
Joel Mandebi Mbongue (University of Florida) and Christophe Bobda (University of Florida)

Accelerating MPI Collectives with FPGAs in the Network and Novel Communicator Support .215
Qingqing Xiong (Boston University), Chen Yang (Boston University), Pouya Haghi (Boston University), Anthony Skjellum (University of Tennessee at Chattanooga), and Martin Herbordt (Boston University)

MeXT-SE: A System-Level Design Tool to Transparently Generate Secure MPSoC .216
Md Jubaer Hossain Pantho (University of Florida) and Christophe Bobda (University of Florida)

Poster Session 3: Tools

Early-Stage Automated Identification Tool for Shared Accelerators .217
Parnian Mokri (Tufts University) and Mark Hempstead (Tufts University)

An Analytical Model of Memory-Bound Applications Compiled with High Level Synthesis .218
Maria A. Dávila-Guzmán (Universidad de Zaragoza), Rubén Gran Tejero (Universidad de Zaragoza), María Villarroya-Gaudó (Universidad de Zaragoza), and Darío Suárez Gracia (Universidad de Zaragoza)

FPGA Virtualization for Deprecated Devices .219
Ian D. Taras (University of Toronto) and Andrew G. Schmidt (University of Southern California)

ZRLMPI: A Unified Programming Model for Reconfigurable Heterogeneous Computing Clusters .220
Burkhard Ringlein (IBM Research Europe; Friedrich-Alexander University Erlangen-Nürnberg), Francois Abel (IBM Research Europe), Alexander Ditter (Friedrich-Alexander University Erlangen-Nürnberg), Beat Weiss (IBM Research Europe), Christoph Hagleitner (IBM Research Europe), and Dietmar Fey (Friedrich-Alexander University Erlangen-Nürnberg)

Designing Domain Specific Computing Systems .221
Anthony M. Cabrera (Washington University in St. Louis) and Roger D. Chamberlain (Washington University in St. Louis)

Poster Session 4: Applications and Architectures

Improving the Availability of Secure Space Links Through the Partial Reconfiguration of FPGAs .222
Emmanuel Lesser (European Space Agency)

An FPGA-Optimized Architecture of Real-Time Farneback Optical Flow .223
Zhe Pan (Zhejiang University), Yuruo Jin (Zhejiang University), Xiaohong Jiang (Zhejiang University), and Jian Wu (Zhejiang University)

High-Performance Parallel Radix Sort on FPGA .224
Bashar Romanous (University of California, Riverside), Mohammadreza Rezvani (University of California, Riverside), Junjie Huang (University of California, Riverside), Daniel Wong (University of California, Riverside), Evangelos E. Papalexakis (University of California, Riverside), Vassilis J. Tsotras (University of California, Riverside), and Walid Najjar (University of California, Riverside)

FPGA-Based Gesture Recognition with Capacitive Sensor Array Using Recurrent Neural Networks .225
Haoyan Liu (University of Arkansas), Atiyehsadat Panahi (University of Arkansas), David Andrews (University of Arkansas), and Alexander Nelson (University of Arkansas)

Gbit/s Non-Binary LDPC Decoders: High-Throughput using High-Level Specifications .226
Oscar Ferraz (University of Coimbra), Srinivasan Subramaniyan (Amrita Vishwa Vidyapeetham), Guohui Wang (Rice University), Joseph R. Cavallaro (Rice University), Gabriel Falcao (University of Coimbra), and Madhura Purnaprajna (Amrita Vishwa Vidyapeetham)

A Quaternary FPGA Architecture Using Floating Gate Memories .227
Ayokunle Fadamiro (Auburn University), Pouyan Rezaie (Auburn University), Christopher Harris (Auburn University), and Spencer Millican (Auburn University)

Rotary Register File: A Micro-Architectural Primitive on FPGA .228
Reza Nakhjavani (University of Toronto) and Jianwen Zhu (University of Toronto)

Poster Session 5: Machine Learning 1

Tiny On-chip Memory Realization of Weight Sparseness Split-CNNs on Low-End FPGAs .229
Akira Jinguji (Tokyo Institute of Technology), Shimpei Sato (Tokyo Institute of Technology), and Hiroki Nakahara (Tokyo Institute of Technology)

An Efficient FPGA-Based Architecture for Contractive Autoencoders .230
Madis Kerner (Tallinn University of Technology), Kalle Tammemäe (Tallinn University of Technology), Jaan Raik (Tallinn University of Technology), and Thomas Hollstein (Tallinn University of Technology; Frankfurt University of Applied Sciences, Frankfurt, Germany)

Systolic-CNN: An OpenCL-Defined Scalable Run-Time-Flexible FPGA Accelerator Architecture for Accelerating Convolutional Neural Network Inference in Cloud/Edge Computing .231
Akshay Dua (Arizona State University), Yixing Li (Arizona State University), and Fengbo Ren (Arizona State University)

Explore Efficient LUT-Based Architecture for Quantized Convolutional Neural Networks on FPGA .232
Yanpeng Cao (Southeast University), Chengcheng Wang (Southeast University), and Yongming Tang (Southeast University)

Realization of Quantized Neural Network for Super-Resolution on PYNQ .233
Feng Yu (Southeast University), Yanpeng Cao (Southeast University), and Yongming Tang (Southeast University)

Scalable Full Hardware Logic Architecture for Gradient Boosted Tree Training .234
Tamon Sadasue (RICOH Company) and Tsuyoshi Isshiki (Tokyo Institute of Technology)

Optimized Distribution of an Accelerated Convolutional Neural Network Across Multiple FPGAs .235
Alaa Maarouf (American University of Beirut), Nour El Droubi (American University of Beirut), Raghid Morcel (American University of Beirut), Hazem Hajj (American University of Beirut), Mazen A. R. Saghir (American University of Beirut), and Haitham Akkary (American University of Beirut)

Poster Session 6: Machine Learning 2

SqueezeJet-3: An Accelerator Utilizing FPGA MPSoCs for Edge CNN Applications .236
Panagiotis Mousouliotis (Aristotle University of Thessaloniki), Ioannis Papaefstathiou (Aristotle University of Thessaloniki), and Loukas Petrou (Aristotle University of Thessaloniki)

Automatic Generation of FPGA Kernels from Open Format CNN Models .237
Dimitrios Danopoulos (NTUA), Christoforos Kachris (Democritus University of Thrace), and Dimitrios Soudris (NTUA)

High-Throughput DNN Inference with LogicNets .238
Yaman Umuroglu (Xilinx Research Labs), Yash Akhauri (Xilinx Research Labs), Nicholas J. Fraser (Xilinx Research Labs), and Michaela Blott (Xilinx Research Labs)

AIgean: An Open Framework for Machine Learning on Heterogeneous Clusters .239
Naif Tarafdar (University of Toronto), Giuseppe Di Guglielmo (Columbia University), Philip C Harris (Massachusetts Institute of Technology), Jeffrey D Krupa (Massachusetts Institute of Technology), Vladimir Loncar (CERN), Dylan S Rankin (Massachusetts Institute of Technology), Nhan Tran (Fermilab), Zhenbin Wu (University of Illinois), Qianfeng Shen (University of Toronto), and Paul Chow (University of Toronto)

FPGA Based High-Throughput Real-Time Feature Extraction for Modulation Classification .240
Joshua Mack (University of Arizona) and Ali Akoglu (University of Arizona)

Accelerating Large Scale GCN Inference on FPGA .241
Bingyi Zhang (University of Southern California), Hanqing Zeng (University of Southern California), and Viktor Prasanna (University of Southern California)

EASpiNN: Effective Automated Spiking Neural Network Evaluation on FPGA .242
Sathish Panchapakesan (Simon Fraser University), Zhenman Fang (Simon Fraser University), and Nitin Chandrachoodan (Indian Institute of Technology - Madras)

A High-Performance Inference Accelerator Exploiting Patterned Sparsity in CNNs . 243
Ning Li (Tsinghua University), Leibo Liu (Tsinghua University), Shaojun Wei (Tsinghua University), and Shouyi Yin (Tsinghua University)

Author Index 245