2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM 2019)

San Diego, California, USA
28 April – 1 May 2019

IEEE

IEEE Catalog Number: CFP19054-POD
ISBN: 978-1-7281-1132-2

Copyright © 2019 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:	CFP19054-POD
ISBN (Print-On-Demand):	978-1-7281-1132-2
ISBN (Online):	978-1-7281-1131-5
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

2019 IEEE 27th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)

FCCM 2019

Table of Contents

Welcome Message from the General and Program Chairs ..... **xiv**
Organizing Committee ..... **xvi**
Program Committee ..... **xvii**
Additional Reviewers ..... **xix**
Sponsors ..... **xx**

Session 1: Open Source Tools

Yosys+nextpnr: An Open Source Framework from Verilog to Bitstream for Commercial FPGAs 1
David Shah (Symbiotic EDA; Imperial College London), Eddie Hung (University of British Columbia; Symbiotic EDA), Clifford Wolf (Symbiotic EDA), Serge Bazanski (Symbiotic EDA), Dan Gisselquist (Symbiotic EDA), and Miodrag Milanovic (Symbiotic EDA)

EFCAD — An Embedded FPGA CAD Tool Flow for Enabling On-chip Self-Compilation 5
Khoa Dang Pham (The University of Manchester), Malte Vesper (The University of Manchester), Dirk Koch (The University of Manchester), and Eddie Hung (University of British Columbia)

Maverick: A Stand-Alone CAD Flow for Partially Reconfigurable FPGA Modules 9
Dallon Glick (Brigham Young University), Jesse Grigg (Brigham Young University), Brent Nelson (Brigham Young University), and Michael Wirthlin (Brigham Young University)

Session 2: Machine Learning 1

An Efficient Hardware Accelerator for Sparse Convolutional Neural Networks on FPGAs 17
Liqiang Lu (Peking University), Jiaming Xie (Peking University), Ruirui Huang (Alibaba Group), Jiansong Zhang (Alibaba Group), Wei Lin (Alibaba Group), and Yun Liang (Peking University)

LUTNet: Rethinking Inference in FPGA Soft Logic 26
Erwei Wang (Imperial College London), James J. Davis (Imperial College London), Peter Y. K. Cheung (Imperial College London), and George A. Constantinides (Imperial College London)

PIR-DSP: An FPGA DSP Block Architecture for Multi-precision Deep Neural Networks 35
SeyedRamin Rasoulinezhad (The University of Sydney), Hao Zhou (Fudan University), Lingli Wang (Fudan University), and Philip H.W. Leong (The University of Sydney)

Towards Efficient Deep Neural Network Training by FPGA-Based Batch-Level Parallelism 45
Cheng Luo (Fudan University), Man-Kit Sit (Imperial College London), Hongxiang Fan (Imperial College London), Shuanglong Liu (Imperial College London), Wayne Luk (Imperial College London), and Ce Guo (Imperial College London)

Session 3: Tools

CRoute: A Fast High-Quality Timing-Driven Connection-Based FPGA Router 53
Dries Vercruyce (Ghent University), Elias Vansteenkiste (Ghent University), and Dirk Stroobandt (Ghent University)

RapidRoute: Fast Assembly of Communication Structures for FPGA Overlays 61
Leo Liu (University of Waterloo), Jay Weng (University of Waterloo), and Nachiket Kapre (University of Waterloo)

Generic Connectivity-Based CGRA Mapping via Integer Linear Programming 65
Matthew J. P. Walker (University of Toronto) and Jason H. Anderson (University of Toronto)

LAMDA: Learning-Assisted Multi-stage Autotuning for FPGA Design Closure 74
Ecenur Ustun (Cornell University), Shaojie Xiang (Cornell University), Jinny Gui (Cornell University), Cunxi Yu (Cornell University), and Zhiru Zhang (Cornell University)

Memory Mapping for Multi-die FPGAs 78
Nils Voss (Imperial College London; Maxeler Technologies), Pablo Quintana (Maxeler Technologies), Oskar Mencer (Maxeler Technologies), Wayne Luk (Imperial College London), and Georgi Gaydadjiev (Imperial College London; Maxeler Technologies)

Impact of FPGA Architecture on Area and Performance of CGRA Overlays 87
Ian Taras (University of Toronto) and Jason H. Anderson (University of Toronto)

An FPGA-Based BWT Accelerator for Bzip2 Data Compression 96
Weikang Qiao (UCLA), Zhenman Fang (Simon Fraser University), Mau-Chung Frank Chang (UCLA), and Jason Cong (UCLA)

Session 4: Applications 1

π-BA: Bundle Adjustment Acceleration on Embedded FPGAs with Co-observation Optimization 100
Shuzhen Qin (Tianjin University), Qiang Liu (Tianjin University), Bo Yu (PerceptIn), and Shaoshan Liu (PerceptIn)

Deep Packet Inspection in FPGAs via Approximate Nondeterministic Automata .109
Milan Ceška (Brno University of Technology), Vojtech Havlena (Brno University of Technology), Lukáš Holík (Brno University of Technology), Jan Korenek (Brno University of Technology), Ondrej Lengál (Brno University of Technology), Denis Matoušek (Brno University of Technology), Jirí Matoušek (Brno University of Technology), Jakub Semric (Brno University of Technology), and Tomáš Vojnar (Brno University of Technology)

Active Stereo Vision with High Resolution on an FPGA .118
Marc Pfeifer (University of Freiburg), Philipp M. Scholl (University of Freiburg), Rainer Voigt (Ensenso GmbH), and Bernd Becker (University of Freiburg)

Hardware Acceleration of Long Read Pairwise Overlapping in Genome Sequencing: A Race Between FPGA and GPU .127
Licheng Guo (University of California, Los Angeles), Jason Lau (University of California, Los Angeles), Zhenyuan Ruan (University of California, Los Angeles), Peng Wei (University of California, Los Angeles), and Jason Cong (University of California, Los Angeles)

Processor Assisted Worklist Scheduling for FPGA Accelerated Graph Processing on a Shared-Memory Platform .136
Yu Wang (Carnegie Mellon University), James C. Hoe (Carnegie Mellon University), and Eriko Nurvitadhi (Intel Labs)

Session 5: Simulation and Infrastructure

MEG: A RISCV-Based System Simulation Infrastructure for Exploring Memory Optimization Using FPGAs and Hybrid Memory Cube .145
Jialiang Zhang (University of Wisconsin-Madison), Yang Liu (University of Wisconsin-Madison), Gaurav Jain (University of Wisconsin-Madison), Yue Zha (University of Wisconsin-Madison), Jonathan Ta (University of Wisconsin-Madison), and Jing Li (University of Wisconsin-Madison)

Enhancing Butterfly Fat Tree NoCs for FPGAs with Lightweight Flow Control .154
Gurshaant Singh Malik (University of Waterloo) and Nachiket Kapre (University of Waterloo)

SimAcc: A Configurable Cycle-Accurate Simulator for Customized Accelerators on CPU-FPGAs SoCs .163
Konstantinos Iordanou (The University of Manchester), Oscar Palomar (The University of Manchester), John Mawer (The University of Manchester), Cosmin Gorgovan (The University of Manchester), Andy Nisbet (The University of Manchester), and Mikel Luján (The University of Manchester)

Session 6: Machine Learning 2

Towards Efficient and Scalable Acceleration of Online Decision Tree Learning on FPGA .172
Zhe Lin (Hong Kong University of Science and Technology), Sharad Sinha (Indian Institute of Technology (IIT) Goa), and Wei Zhang (Hong Kong University of Science and Technology)

T2S-Tensor: Productively Generating High-Performance Spatial Hardware for Dense Tensor Computations .181
Nitish Srivastava (Cornell University), Hongbo Rong (Intel), Prithayan Barua (Georgia Institute of Technology), Guanyu Feng (Tsinghua University), Huanqi Cao (Tsinghua University), Zhiru Zhang (Cornell University), David Albonesi (Cornell University), Vivek Sarkar (Georgia Institute of Technology), Wenguang Chen (Tsinghua University), Paul Petersen (Intel), Geoff Lowney (Intel), Adam Herr (Intel), Christopher Hughes (Intel), Timothy Mattson (Intel), and Pradeep Dubey (Intel)

SparseHD: Algorithm-Hardware Co-optimization for Efficient High-Dimensional Computing .190
Mohsen Imani (University of California, San Diego), Sahand Salamat (University of California, San Diego), Behnam Khaleghi (University of California, San Diego), Mohammad Samragh (University of California, San Diego), Farinaz Koushanfar (University of California, San Diego), and Tajana Rosing (University of California, San Diego)

Why Compete When You Can Work Together: FPGA-ASIC Integration for Persistent RNNs .199
Eriko Nurvitadhi (Intel Corporation), Dongup Kwon (Intel Corporation), Ali Jafari (Intel Corporation), Andrew Boutros (Intel Corporation), Jaewoong Sim (Intel Corporation), Phillip Tomson (Intel Corporation), Huseyin Sumbul (Intel Corporation), Gregory Chen (Intel Corporation), Phil Knag (Intel Corporation), Raghavan Kumar (Intel Corporation), Ram Krishnamurthy (Intel Corporation), Sergey Gribok (Intel Corporation), Bogdan Pasca (Intel Corporation), Martin Langhammer (Intel Corporation), Debbie Marr (Intel Corporation), and Aravind Dasu (Intel Corporation)

Session 7: Programming

Design Patterns for Code Reuse in HLS Packet Processing Pipelines .208
Haggai Eran (Technion), Lior Zeno (Technion), Zsolt István (IMDEA Software Institute), and Mark Silberstein (Technion)

Module-per-Object: A Human-Driven Methodology for C++-Based High-Level Synthesis Design .218
Jeferson Santiago da Silva (Polytechnique Montréal), François-Raymond Boyer (Polytechnique Montréal), and J.M. Pierre Langlois (Polytechnique Montréal)

Templatised Soft Floating-Point for High-Level Synthesis .227
David B. Thomas (Imperial College London)

Exploiting Irregular Memory Parallelism in Quasi-Stencils through Nonlinear Transformation .236
Juan Escobedo (University of Central Florida) and Mingjie Lin (University of Central Florida)

Session 8: Applications 2

FP-AMR: A Reconfigurable Fabric Framework for Adaptive Mesh Refinement Applications .245
Tianqi Wang (University of Science and Technology of China), Tong Geng (Boston University), Xi Jin (University of Science and Technology of China), and Martin Herbordt (Boston University)

Compressed Sensing MRI Reconstruction on Intel HARPv2 .254
Yushan Su (Princeton University), Michael Anderson (Intel Labs), Jonathan I. Tamir (University of California, Berkeley), Michael Lustig (University of California, Berkeley), and Kai Li (Princeton University)

GhostSZ: A Transparent FPGA-Accelerated Lossy Compression Framework .258
Qingqing Xiong (Boston University), Rushi Patel (Boston University), Chen Yang (Boston University), Tong Geng (Boston University), Anthony Skjellum (University of Tennessee at Chattanooga), and Martin C. Herbordt (Boston University)

Efficient Hardware Acceleration for Design Diversity Calculation to Mitigate Common Mode Failures .267
Maheshwaran Ramesh Babu (The University of Texas at Dallas), Farah Naz Taher (The University of Texas at Dallas), Anjana Balachandran (The Hong Kong Polytechnic University), and Benjamin Carrion Schafer (The University of Texas at Dallas)

Session 9: Security and Reliability

Fast Voltage Transients on FPGAs: Impact and Mitigation Strategies .271
Linda L. Shen (University of Toronto), Ibrahim Ahmed (University of Toronto), and Vaughn Betz (University of Toronto)

FASE: FPGA Acceleration of Secure Function Evaluation .280
Siam U. Hussain (University of California, San Diego) and Farinaz Koushanfar (University of California, San Diego)

Session 10: Arithmetic

Rethinking Integer Divider Design for FPGA-Based Soft-Processors .289
Eric Matthews (Simon Fraser University), Alec Lu (Simon Fraser University), Zhenman Fang (Simon Fraser University), and Lesley Shannon (Simon Fraser University)

High Precision, High Performance FPGA Adders .298
Martin Langhammer (Intel Corporation), Bogdan Pasca (Intel Corporation), and Gregg Baeckler (Intel Corporation)

Poster Session 1: Tools

Automated Tool and Runtime Support for Fine-Grain Reconfiguration in Highly Flexible Reconfigurable Systems .307
Rafael Zamacola (Universidad Politécnica de Madrid), Alberto García Martínez (Universidad Politécnica de Madrid), Javier Mora (Universidad Politécnica de Madrid), Andrés Otero (Universidad Politécnica de Madrid), and Eduardo de la Torre (Universidad Politécnica de Madrid)

AutoPhase: Compiler Phase-Ordering for HLS with Deep Reinforcement Learning .308
Qijing Huang (University of California, Berkeley), Ameer Haj-Ali (University of California, Berkeley), William Moses (Massachusetts Institute of Technology), John Xiang (University of California, Berkeley), Ion Stoica (University of California, Berkeley), Krste Asanovic (University of California, Berkeley), and John Wawrzynek (University of California, Berkeley)

Efficient FPGA Floorplanning for Partial Reconfiguration-Based Applications .309
Norbert Deak (National Instruments), Octavian Cret (Technical University of Cluj-Napoca), and Horia Hedesiu (Technical University of Cluj-Napoca)

Towards Prototyping and Acceleration of Java Programs onto Intel FPGAs .310
Michail Papadimitriou (The University of Manchester), Juan Fumero (The University of Manchester), Athanasios Stratikopoulos (The University of Manchester), and Christos Kotselidis (The University of Manchester)

Sonar: Writing Testbenches through Python .311
Varun Sharma (University of Toronto), Naif Tarafdar (University of Toronto), and Paul Chow (University of Toronto)

Raparo: Resource-Level Angle-Based Parallel Routing for FPGAs .312
Minghua Shen (Sun Yat-Sen University) and Nong Xiao (Sun Yat-Sen University)

Automated Acceleration of Dataflow-Oriented C Applications on FPGA-Based Systems .313
Francesco Peverelli (Politecnico di Milano), Marco Rabozzi (Politecnico di Milano), Salvatore Cardamone (University Chemical Laboratory), Emanuele Del Sozzo (Politecnico di Milano), Alex J. W. Thom (University Chemical Laboratory), Marco D. Santambrogio (Politecnico di Milano), and Lorenzo Di Tucci (Politecnico di Milano)

Automated Design Space Exploration and Roofline Analysis for FPGA-Based HLS Applications .314
Marco Siracusa (Politecnico di Milano), Marco Rabozzi (Politecnico di Milano), Emanuele Del Sozzo (Politecnico di Milano), Marco D. Santambrogio (Politecnico di Milano), and Lorenzo Di Tucci (Politecnico di Milano)

Formalizing Loop-Carried Dependencies in Coq for High-Level Synthesis .315
Florian Faissole (Inria), George A. Constantinides (Imperial College London), and David Thomas (Imperial College London)

Poster Session 2: Neural Networks and Vision

Exploring the Random Network of Hodgkin and Huxley Neurons with Exponential Synaptic Conductances on OpenCL FPGA Platform .316
Zheming Jin (Argonne National Laboratory) and Hal Finkel (Argonne National Laboratory)

A Scalable OpenCL-Based FPGA Accelerator for YOLOv2 .317
Ke Xu (Beijing Jiaotong University), Xiaoyun Wang (Beijing Jiaotong University), and Dong Wang (Beijing Key Laboratory of Advanced Information Science and Network Technology)

Model-Extraction Attack Against FPGA-DNN Accelerator Utilizing Correlation Electromagnetic Analysis .318
Kota Yoshida (Ritsumeikan University), Takaya Kubota (Ritsumeikan University), Mitsuru Shiozaki (Ritsumeikan University), and Takeshi Fujino (Ritsumeikan University)

SimBNN: A Similarity-Aware Binarized Neural Network Acceleration Framework .319
Cheng Fu (University of California, San Diego), Shilin Zhu (University of California, San Diego), Huili Chen (University of California, San Diego), Farinaz Koushanfar (University of California, San Diego), Hao Su (University of California, San Diego), and Jishen Zhao (University of California, San Diego)

KPynq: A Work-Efficient Triangle-Inequality Based K-Means on FPGA .320
Yuke Wang (University of California, Santa Barbara), Zhaorui Zeng (University of California, Santa Barbara), Boyuan Feng (University of California, Santa Barbara), Lei Deng (University of California, Santa Barbara), and Yufei Ding (University of California, Santa Barbara)

A High Throughput and Energy-Efficient Retina-Inspired Tone Mapping Processor .321
Lili Liu (University of Electronic Science and Technology of China), Xiaoqiang Xiang (University of Electronic Science and Technology of China), Yuxiang Xie (University of Electronic Science and Technology of China), Yongjie Li (University of Electronic Science and Technology of China), Bo Yan (University of Electronic Science and Technology of China), and Jun Zhou (University of Electronic Science and Technology of China)

An OpenCL-Based Acceleration for Canny Algorithm Using a Heterogeneous CPU-FPGA Platform .322
Samah Rahamneh (Western Michigan University) and Lina Sawalha (Western Michigan University)

Poster Session 3: Networking and Resource Sharing

Scalable P4 Deparser for Speeds Over 100 Gbps .323
Jakub Cabal (CESNET a.l.e.), Pavel Benacek (CESNET a.l.e.), Jana Foltova (CESNET a.l.e.), and Juraj Holub (Brno University of Technology)

Wire-Speed Multirate Accelerator for Aggregation Operations on Sorted Data .324
Sang-Woo Jun (University of California, Irvine) and Arvind Arvind (Massachusetts Institute of Technology)

Hybrid XML Parser Based on Software and Hardware Co-design .325
Zhe Pan (Zhejiang University), Xiaohong Jiang (Zhejiang University), Jian Wu (Zhejiang University), and Xiang Li (Zhejiang University)

Sorting Large Data Sets with FPGA-Accelerated Samplesort .326
Han Chen (Stony Brook University), Sergey Madaminov (Stony Brook University), Michael Ferdman (Stony Brook University), and Peter Milder (Stony Brook University)

Cost-Effective Energy Monitoring of a Zynq-Based Real-Time System Including Dual Gigabit Ethernet .327
Martin Geier (Technical University of Munich), Dominik Faller (Technical University of Munich), Marian Brändle (Technical University of Munich), and Samarjit Chakraborty (Technical University of Munich)

Improved Techniques for Sensing Intra-Device Side Channel Leakage .328
William Hunter (Georgia Institute of Technology), Christopher McCarty (Georgia Institute of Technology), and Lee Lerner (Georgia Institute of Technology)

Safe Task Interruption for FPGAs .329
Sameh Attia (University of Toronto) and Vaughn Betz (University of Toronto)

Poster Session 4: Applications

OpenCL Kernel Vectorization on the CPU, GPU, and FPGA: A Case Study with Frequent Pattern Compression .330
Zheming Jin (Argonne National Laboratory) and Hal Finkel (Argonne National Laboratory)

A 4.8x Faster FPGA-Based Iterative Closest Point Accelerator for Object Pose Estimation of Picking Robot Applications .331
Atsutake Kosuge (Hitachi Ltd.), Keisuke Yamamoto (Hitachi Ltd.), Yukinori Akamine (Hitachi Ltd.), Taizo Yamawaki (Hitachi Ltd.), and Takashi Oshima (Hitachi Ltd.)

Monobit Wideband Receiver with Integrated Dithering in FPGA .332
Dan Pritsker (Intel Corporation) and Colman Cheung (Intel Corporation)

An FPGA-Based Computing Infrastructure Tailored to Efficiently Scaffold Genome Sequences .333
Alberto Zeni (Politecnico di Milano), Matteo Crespi (Politecnico di Milano), Lorenzo Di Tucci (Politecnico di Milano), and Marco D. Santambrogio (Politecnico di Milano)

FlexGibbs: Reconfigurable Parallel Gibbs Sampling Accelerator for Structured Graphs .334
Glenn G. Ko (Harvard University), Yuji Chai (University of Illinois at Urbana-Champaign), Rob A. Rutenbar (University of Pittsburgh), David Brooks (Harvard University), and Gu-Yeon Wei (Harvard University)

A Fine-Grained Parallel Snappy Decompressor for FPGAs Using a Relaxed Execution Model .335
Jian Fang (Delft University of Technology), Jianyu Chen (Delft University of Technology), Jinho Lee (IBM Research), Zaid Al-Ars (Delft University of Technology), and H. Peter Hofstee (Delft University of Technology; IBM Research)

Analyzing the Energy-Efficiency of Vision Kernels on Embedded CPU, GPU and FPGA Platforms .336
Murad Qasaimeh (Iowa State University), Joseph Zambreno (Iowa State University), Phillip H. Jones (Iowa State University), Kristof Denolf (Xilinx Research Labs), Jack Lo (Xilinx Research Labs), and Kees Vissers (Xilinx Research Labs)

Large-Scale and High-Throughput QR Decomposition on an FPGA .337
Dajung Lee (Intel Corporation), Andrei Hagiescu (Intel Corporation), and Dan Pritsker (Intel Corporation)

Flexi-AES: A Highly-Parameterizable Cipher for a Wide Range of Design Constraints 338
Sergiu Mosanu (University of Virginia), Xinfei Guo (University of Virginia), Mohamed El-Hadedy (California State Polytechnic University, Pomona), Lorena Anghel (Grenoble Institute of Engineering), and Mircea Stan (University of Virginia)

Author Index 339