Poster Session 1 (Chair: George Constantinides) Monday February 25, 10:15

Scheduling Data in Neural Network Applications
Thaddeus Koehn(2), Peter Athanas(3, 1)
State University(1), Northrup Grumman(2), Virginia Polytechnic Institute(3)

Fault Testing a Synthesizable Embedded Processor at Gate Level using UltraScale FPGA Emulation
Tom Mannos(2), Brian Dziki(1), Moslema Sharif(2)
Department of Defense(1), Sandia National Laboratories(2)

Base64 Encoding on OpenCL FPGA Platform
Zheming Jin, Hal Finkel: Argonne National Laboratory

Scalable High Performance SDN Switch Architecture on FPGA for Core Networks
Sasindu Wijeratne, Ashen Ekanayake, Sandaruwan Jayaweera, Danuka Ravishan, Ajith Pasqual
University of Moratuwa

A Deep-Reinforcement-Learning-Based Scheduler for High-Level Synthesis
Hongzheng Chen, Minghua Shen: Sun Yat-sen University

Accelerating 3D CNN-based Lung Nodule Segmentation on a Multi-FPGA System
Junzhong Shen, Deguang Wang, You Huang, Mei Wen, Chunyuan Zhang: National University of Defense Technology

SparseBNN: Joint Algorithm/Hardware Optimization to Exploit Structured Sparsity in Binary Neural Network
Xin He, Liu Ke, Xuan Zhang: Washington University in St. Louis

A Deep Learning Inference Accelerator Based on Model Compression on FPGA
Lu Jing, Jun Liu, FuHai Yu: Inspur Corporation

Sparse Winograd Convolutional Neural Networks on Small-scale Systolic Arrays
Feng Shi(1), Haochen Li(1), Yuhe Gao(2), Benjamin Kuschner(1), Song-Chun Zhu(1)
University of California, Los Angeles(1), University of Hong Kong(2)

HSC-FPGA: A Hybrid Spin/Charge FPGA Leveraging the Cooperating Strengths of CMOS and MTJ Devices
Ramtin Zand, Ronald DeMara: University of Central Florida

Evaluating and Enhancing Intel Stratix 10 FPGAs for Persistent Real-Time AI
Eriko Nurvitadhi, Dongup Kwon, Ali Jafari, Andrew Boutros, Jaewoong Sim, Phillip Tomson,
Huseyin Sumbul, Gregory Chen, Phil Knag, Raghavan Kumar, Ram Krishnamurthy,
Debbie Marr, Sergey Gribok, Bogdan Pasca, Martin Langhammer, Aravind Dasu
Intel Corporation

Parrot: A More Effective Parallel Routing Approach to FPGAs
Minghua Shen, Nong Xiao: Sun Yat-Sen University

A Reconfigurable Accelerator for Sparse Convolutional Neural Networks
Weijie You, Chang Wu: Fudan University

Overcoming Data Transfer Bottlenecks in DNN Accelerators via Layer-Conscious Memory Managment
Xuechao Wei(3, 2), Yun Liang(3), Peng Zhang(2), Cody Yu(1), Jason Cong(1, 2)
Peking University, UCLA (1), Falcon Computing Solutions, Inc.(2), Peking University(3)

A Pixel-Parallel Virtual-Image Architecture for High Performance and Power Efficient Graph Cuts Inference
Tianqi Gao(1), Rob Rutenbar(2)
University of Illinois at Urbana Champaign(1), University of Pittsburgh(2)

Unleashing the Power of Soft Logic for Convolutional Neural Network Acceleration via Product Quantization
Jialiang Zhang, Jing Li: University of Wisconsin-Madison

Highly Efficient Sparse Neural Network Computing: Hardware and Software Solutions
Yanjie Gu(2), Jian Yu(2), Tieli Sun(2), Chen Pan(2), Zhenhao Feng(2), Liewei Xu(2), Chang Wu(1)
Fudan University(1), Shanghai Fudan Microelectronics Group Company Limited(2)

FPGA-based Distributed Edge Training of SVM
Jyotikrishna Dass, Yashwardhan Narawane, Rabi Mahapatra, Vivek Sarin: Texas A&M University

Transistor-Level Optimization Methodology for GRM FPGA Interconnect Circuits
Zhengjie Li, Yuanlong Xiao, Yufan Zhang, Yunbing Pang, Jian Wang, Jinmei Lai: Fudan University

Poster Session 2 (Chair: Phillip Leong) Monday February 25, 15:10

PAI-FCNN: FPGA Based CNN Inference System

Lansong Diao, Zhao Jiang, Hao Liang, Chang'an Ye, Kai Chen, Li Ding, Shunli Dou, Meng Sun, Lixue Xia, Jiansong Zhang, Wei Lin
Alibaba Group

JuxtaPiton: Enabling Heterogeneous-ISA Research with RISC-V and SPARC FPGA Soft-cores

Katie Lim(1), Jonathan Balkind(2), David Wentzlaff(2): University of Washington(1), Princeton University(2)

MODA-PSO: Towards Fast Hard Block Legalization for Analytical FPGA Placement

Yun Zhou, Dries Vercruyce, Dirk Stroobandt: Ghent University

A PYNQ-compliant Online Platform for Zynq-based DNN Developers

Chen Chen(2), Jun Xia(2), Wenmin Yang(1), Kang Li(2), Zhilei Chai(2): OpenHEC Lab(1), Jiangnan University(2)

SwitchAgg: A Further Step Towards In-Network Computation

Fan Yang, Zhan Wang, Xiaoxiao Ma, Guojun Yuan, Xuejun An: Institute of Computing Technology, Chinese Academy of Sciences

A Fine-Grained Sparse Accelerator for Multi-Precision DNN

Shulin Zeng(1), Yujun Lin(3), Shuang Liang(1), Junlong Kang(2), Dongliang Xie(2), Yi Shan(2), Song Han(3), Yu Wang(1), Huazhong Yang(1)
Tsinghua University(1), Xilinx(2), Massachusetts Institute of Technology(3)

Building FPGA State Machines from Sequential Code

Carl-Johannes Johnsen, Kenneth Skovhede: University of Copenhagen

Design and Implementation of a Deterministic FPGA Router on a CPU+FPGA Acceleration Platform

Dario Korolija, Mirjana Stojilovic: EPFL

An FPGA-based Fine Tuning Accelerator for a Sparse CNN

Hiroki Nakahara, Akira Jinguji, Masayuki Shimoda, Shimpei Sato: Tokyo Institute of Technology

Embracing Systolic: Super Systolization of Large-Scale Circulant Matrix-vector Multiplication on FPGA with Subquadratic Space Complexity

Jiafeng Xie(1), Chiou-Yng Lee(2): Villanova University(1), Lunghwa University of Science and Technology(2)

Dataflow Systolic Array Implementations of Matrix Decomposition Using High Level Synthesis

Jie Liu(1), Jason Cong(2): Tsinghua University(1), UCLA(2)

Speedy: An Accelerator for Sparse Convolutional Neural Networks on FPGAs

Liqiang Lu(2), Yun Liang(2), Ruirui Huang(1), Wei Lin(1), Xiaoyuan Cui(1), Jiansong Zhang(1)
Alibaba Group(1), Peking University(2)

DNNVM : End-to-End Compiler Leveraging Operation Fusion on FPGA-based CNN Accelerators

Yu Xing(1), Shuang Liang(2), Lingzhi Sui(1), Zhen Zhang(1), Jiantao Qiu(2), Xijie Jia(1), Xin Liu(1), Yushun Wang(1), Yi Shan(1), Yu Wang(2)
Xilinx(1) Tsinghua University(2)

A Hybrid Data-Consistent Framework for Link-Aware AccessManagement in Emerging CPU-FPGA Platforms

Liang Feng(1), Jieru Zhao(1), Tingyuan Liang(1), Sharad Sinha(2), Wei Zhang(1)
Hong Kong University of Science and Technology(1), Indian Institute of Technology Goa(2)

On Feasibility of FPGAs Without Dedicated Programmable Interconnect Structure

Anastasiia Kucherenko, Stefan Nikolic, Paolo Ienne: EPFL

Fast Confidence Detection: One Hot Way to Detect Adversarial Attacks via Sensor Pattern Noise Fingerprinting

Yazhu Lan(1), Qingli Guo(4), Guohe Zhang(2), Yuanchao Xu(3), Kent Nixon(1), Hai Li(1), Yiran Chen(1)
Duke University(1), Xian Jiaotong University(2), Capital Normal University(3), University of Chinese Academy of Sciences(4)

FTConv: FPGA Acceleration for Transposed Convolution Layers in Deep Neural Networks

Zhucheng Tang, Guojie Luo, Ming Jiang: Peking University

Compressed CNN Training with FPGA-based Accelerator

Kaiyuan Guo, Shuang Liang, Jincheng Yu, Xuefei Ning, Wenshuo Li, Yu Wang, Huazhong Yang: Tsinghua University

Optimizing Order-Associative Kernel Computation with Joint Memory Banking and Data Reuse

Juan Escobedo, Mingjie Lin: University of Central Florida

PVT-Aware Sensing and Voltage Scaling for Energy Efficient FPGAs

Konstantinos Maragos(2), George Lentaris(2), Dimitrios Soudris(2), Vasilis F. Pavlidis(1)
The University of Manchester(1), National Technical University of Athens(2)

Software Hardware Co-Optimized BFS on FPGAs

Zachary Sherer, Eric Finnerty, Yan Luo, Hang Liu: University of Massachusetts Lowell

Poster Session 3 (Chair: Kees Vissers) Tuesday February 26, 15:00

Nuclear Reactor Simulations on OpenCL FPGA Platform

Zheming Jin, Hal Finkel: Argonne National Laboratory

Storage Mirroring for Bare-Metal Systems on FPGA Devices

Dan Turicu(1), Octavian Cret(2), Lucia Vacariu(2)
Bitdefender, Technical University of Cluj-Napoca(1), Technical University of Cluj-Napoca(2)

Fast Inference of Deep Neural Networks for Real-time Particle Physics Applications

Javier Duarte(2), Song Han(1), Philip Harris(1), Sergo Jindariani(2), Edward Kreinar(6), Benjamin Kreis(2), Vladimir Loncar(5),
Jennifer Ngadiuba(5), Maurizio Pierini(5), Dylan Rankin(1), Ryan Rivera(2), Sioni Summers(4), Nhan Tran(2), Zhenbin Wu(3)
Massachusetts Institute of Technology(1), Fermilab(2), University of Illinois Chicago(3), Imperial College(4), CERN(5), Hawkeye 360(6)

XFER: A Novel Design to Achieve Super-Linear Performance on Multiple FPGAs for Real-Time AI

Weiwen Jiang(4, 5), Xinyi Zhang(4), Edwin Sha(2), Qingfeng Zhuge(2), Lei Yang(5, 3), Yiyu Shi(1), Jingtong Hu(4)
University of Notre Dame(1), East China Normal University(2), UC Irvine(3), University of Pittsburgh(4), Chongqing University(5)

An Energy-Efficient FPGA Implementation of an LSTM Network Using Approximate Computing

Elham Azari, Aykut Dengi, Sarma Vrudhula: Arizona State University

Towards Fast and Energy-Efficient Binarized Neural Network Inference on FPGA

Cheng Fu(2, 1), Shilin Zhu(1), Hao Su(1), Ching-En Lee(2), Jishen Zhao(1)
University of California San Diego (1), Iluvatar CoreX(2)

BRISC-V: An Open-Source Architecture Design Space Exploration Toolbox

Sahan Bandara, Alan Ehret, Donato Kava, Michel Kinsy: Boston University

Maverick: A Stand-alone CAD Flow for Xilinx 7-Series FPGAs

Dallon Glick, Jesse Grigg, Brent Nelson, Michael Wirthlin
NSF Center for Space, High-Performance, and Resilient Computing (SHREC) & Brigham Young University

Hierarchical FPGA Fabrics using 2D-Benes-BFT-Pyramid Network Layouts with Optimizations

Venkat Konda: Konda Technologies Inc.

Flat FPGA Fabrics Derived from 2D-Benes-BFT-Pyramid Networks with Optimizations and Enhancements

Venkat Konda: Konda Technologies Inc.

HOTMeTaL: Hardware Optimization Tool for Memory Table and Logic Conversion

Michael Kapralos, John Chandy: University of Connecticut

Engaging Heterogeneous FPGAs in the Cloud

Ke Zhang, Yisong Chang, Mingyu Chen, Yungang Bao, Zhiwei Xu: ICT, CAS; University of Chinese Academy of Sciences

Enhancing Butterfly Fat Tree NoCs for FPGAs with Lightweight Flow Control

Gurshaant Singh Malik, Nachiket Kapre: University of Waterloo

Efficient FPGA Implementation of Conjugate Gradient Methods for Laplacian System using HLS

Sahithi Rampalli, Natasha Sehgal, Ishita Bindlish, Tanya Tyagi, Pawan Kumar
International Institute of Information Technology

A FPGA Implementation of Farneback Optical Flow by High-Level Synthesis

Chia-Wei Chang, Zi-Qi Zhong, Jing-Jia Liou: National Tsing Hua University

Efficient Acceleration of CNNs for Semantic Segmentation on FPGAs

Sebastian Vogel(2), Jannik Springer(2, 1), Andre Guntoro(2), Gerd Ascheid(1)
RWTH Aachen University(1), Robert Bosch GmbH(2)