



**Conference Program of the  
25<sup>th</sup> ACM/SIGDA International Symposium on  
Field-Programmable Gate Arrays**

**February 22 - 24, 2017**

**Monterey Marriott Hotel, Monterey, California**

# FPGA 2017 Sponsors and Corporate Patrons

## Sponsors



## Gold Corporate Patrons



## Silver Corporate Patrons



Trimberger Family  
Foundation

Logistics Support:

**Wednesday February 22 (All Technical Sessions in San Carlos 2-4)**

**Int'l Workshop on Overlay Architectures for FPGAs (OLAF)**

Chair: Hayden So, The University of Hong Kong; Co-Chair: John Wawryznek, UC Berkeley

9:00 Welcome and Opening Remarks  
*John Wawryznek (UC Berkeley)*

9:10 Paper Presentations (<http://olaf.eecs.berkeley.edu/program>)

12:00 **Lunch (Ferrantes Room, 10<sup>th</sup> Floor)**

**Afternoon Special Session: The Role of FPGAs in Machine Learning**

Chair: Andrew Ling, Intel

1:30 Deep Learning -- Tutorial and Recent Trends  
*Song Han (Stanford and DeePhi)*

2:30 **Can FPGAs Beat GPUs in Accelerating Next-Generation Deep Neural Networks?**  
*Eriko Nurvitadhi, Ganesh Venkatesh, Jaewoong Sim, Debbie Marr, Randy Huang, Jason Gee Hock Ong, Yeong Tat Liew, Srivatsan Krishnan, Duncan Moss, Suchit Subhaschandra, Guy Boudoukh, Intel*

3:00 Break

**Accelerating Binarized Convolutional Neural Networks with Software-Programmable FPGAs**  
3:30 *Ritchie Zhao<sup>1</sup>, Weinan Song<sup>1</sup>, Wentao Zhang<sup>1</sup>, Tianwei Xing<sup>2</sup>, Jeng-Hau Lin<sup>3</sup>, Mani Srivastava<sup>2</sup>, Rajesh Gupta<sup>3</sup>, Zhiru Zhang<sup>1</sup>*  
<sup>1</sup>Cornell University, <sup>2</sup>UCLA, <sup>3</sup>UCSD

3:55 **Improving the Performance of OpenCL-based FPGA Accelerator for Convolutional Neural Network**  
*Jialiang Zhang and Jing Li, UW-Madison*

4:20 **Frequency Domain Acceleration of Convolutional Neural Networks on CPU-FPGA Shared Memory System**  
*Chi Zhang and Viktor Prasanna, USC*

4:45 **Optimizing Loop Operation and Dataflow in FPGA Acceleration of Deep Convolutional Neural Networks**  
*Yufei Ma, Yu Cao, Sarma Vrudhula, Jae-sun Seo, Arizona State University*

7:00 **Opening Reception (Ferrantes Room, 10<sup>th</sup> Floor)**

## Thursday February 23 (All Technical Sessions in San Carlos 2-4)

### 8:00 Continental Breakfast

8:45 Welcome and Opening Remarks  
*Jason Anderson (University of Toronto), Jonathan Greene (Microsemi)*

### Machine Learning

Chair: Jason Cong, UCLA

### 9:00 An OpenCL Deep Learning Accelerator on Arria 10

*Utku Aydonat, Shane O'Connell, Davor Capalija, Andrew Ling, Gordon Chiu, Intel*

### 9:25 FINN: A Framework for Fast, Scalable Binarized Neural Network Inference

*Yaman Umuroglu<sup>1</sup>, Nicholas J. Fraser<sup>2</sup>, Giulio Gambardella<sup>3</sup>, Michaela Blott<sup>3</sup>, Philip Leong<sup>4</sup>, Magnus Jahre<sup>5</sup>, Kees Vissers<sup>3</sup>*

<sup>1</sup>Xilinx and Norwegian University of Science and Technology, <sup>2</sup>Xilinx and University of Sydney, <sup>3</sup>Xilinx, <sup>4</sup>University of Sydney, <sup>5</sup>Norwegian University of Science and Technology

### 9:50 ESE: Efficient Speech Recognition Engine with Compressed LSTM on FPGA

*Song Han<sup>1</sup>, Junlong Kang<sup>2</sup>, Huizi Mao<sup>1</sup>, Yiming Hu<sup>3</sup>, Xin Li<sup>2</sup>, Yubin Li<sup>2</sup>, Dongliang Xie<sup>2</sup>, Hong Luo<sup>2</sup>, Song Yao<sup>2</sup>, Yu Wang<sup>3</sup>, Huazhong Yang<sup>3</sup>, Bill Dally<sup>1</sup>*

<sup>1</sup>Stanford University, <sup>2</sup>DeePhi, <sup>3</sup>Tsinghua University

### 10:15 Poster Session 1 and Break (San Carlos 1)

### Interconnect and Routing

Chair: Sinan Kaptanoglu, Microsemi

### 11:15 Quality-Time Tradeoffs in Component-Specific Mapping

*Hans Giesen<sup>1</sup>, Raphael Rubin<sup>1</sup>, Benjamin Gojman<sup>2</sup>, Andre DeHon<sup>1</sup>*

<sup>1</sup>University of Pennsylvania, <sup>2</sup>Google

### 11:40 Synchronization Constraints for Interconnect Synthesis

*Alex Rodionov and Jonathan Rose, University of Toronto*

### 12:05 Corolla: GPU-Accelerated FPGA Routing Based on Subgraph Dynamic Expansion

*Minghua Shen and Guojie Luo, Peking University*

### 12:30 Lunch (Ferrantes Room, 10<sup>th</sup> Floor)

### Architecture

Chair: Steve Wilton, University of British Columbia

### 2:00 Don't Forget the Memory: Automatic Block RAM Modelling, Optimization, and Architecture Exploration

*Sadegh Yazdanshenas, Kosuke Tatsumura, Vaughn Betz, University of Toronto*

### 2:25 Automatic Construction of Program-Optimized FPGA Memory Networks

*Hsin-Jung Yang<sup>1</sup>, Kermin Fleming<sup>2</sup>, Felix Winterstein<sup>3</sup>, Annie Chen<sup>1</sup>, Michael Adler<sup>4</sup>, Joel Emer<sup>1</sup>*

<sup>1</sup>MIT, <sup>2</sup>Intel, <sup>3</sup>Imperial College London, <sup>4</sup>Intel Corporation

2:50 **NAND-NOR: A Compact, Fast, and Delay Balanced FPGA Logic Element**  
*Zhihong Huang<sup>1</sup>, Xing Wei<sup>1</sup>, Grace Zgheib<sup>2</sup>, Wei Li<sup>1</sup>, Yu Lin<sup>1</sup>, Zhenghong Jiang<sup>1</sup>, Kaihui Tu<sup>1</sup>, Paolo Ienne<sup>2</sup>, Haigang Yang<sup>1</sup>*  
<sup>1</sup>Chinese Academy of Sciences, <sup>2</sup>EPFL

2:55 **120-core microAptiv MIPS Overlay for the Terasic DE5-NET FPGA board**  
*Nachiket Kapre<sup>1</sup>, Prashanth Ravi<sup>2</sup>, Gourav Modi<sup>3</sup>, Chethan Kumar H B<sup>2</sup>*  
<sup>1</sup>University of Waterloo, <sup>2</sup>Nanyang Technological University, <sup>3</sup>NTU

3:00 **Poster Session 2 and Break (San Carlos 1)**

### **CAD Tools**

Chair: Lesley Shannon, Simon Fraser University

4:00 **A Parallelized Iterative Improvement Approach to Area Optimization for LUT-Based Technology Mapping**  
*Gai Liu and Zhiru Zhang, Cornell University*

4:25 **A Parallel Bandit-Based Approach for Autotuning FPGA Compilation**  
*Chang Xu<sup>1</sup>, Gai Liu<sup>2</sup>, Ritchie Zhao<sup>2</sup>, Stephen Yang<sup>3</sup>, Guojie Luo<sup>1</sup>, Zhiru Zhang<sup>2</sup>*  
<sup>1</sup>Peking University, <sup>2</sup>Cornell University, <sup>3</sup>Xilinx

6:30 **Banquet (San Carlos 2-4)**

7:45 **Banquet Panel: FPGAs in the Cloud**

Chair: George Constantinides, Imperial College London

Panelists: Andrew Putnam (Microsoft, USA), Wei Qi (Baidu, China),

Gaurav Singh (Xilinx, USA), Mark Shand (Waymo, USA),

Ling Shao (IBM Research, China), Richard Veitch (Maxeler, USA)

## **Friday February 24 (All Technical Sessions in San Carlos 2-4)**

8:00 **Continental Breakfast**

### **High-Level Synthesis -- Tools and Applications**

Chair: Stephen Neuendorffer, Xilinx

9:00 **Hardware Synthesis of Weakly Consistent C Concurrency**  
*Nadesh Ramanathan, Shane Fleming, John Wickerson, George Constantinides,*  
Imperial College London

9:25 **A New Approach to Automatic Memory Banking using Trace-Based Address Mining**  
*Yuan Zhou, Khalid Al-Hawaj, Zhiru Zhang, Cornell University*

9:50 **Dynamic Hazard Resolution for Pipelining Irregular Loops in High-Level Synthesis**  
*Steve Dai<sup>1</sup>, Ritchie Zhao<sup>1</sup>, Gai Liu<sup>1</sup>, Shreesha Srinath<sup>1</sup>, Udit Gupta<sup>2</sup>, Christopher Batten<sup>1</sup>, Zhiru Zhang<sup>1</sup>*  
<sup>1</sup>Cornell University, <sup>2</sup>Harvard University

### **Accelerating Face Detection on Programmable SoC Using C-Based Synthesis**

9:55 *Nitish Srivastava<sup>1</sup>, Steve Dai<sup>1</sup>, Rajit Manohar<sup>2</sup>, Zhiru Zhang<sup>1</sup>*

<sup>1</sup>Cornell University, <sup>2</sup>Cornell NYC Tech

### **Packet Matching on FPGAs Using HMC Memory: Towards One Million Rules**

10:00 *Daniel Rozhko, Geoffrey Elliott, Daniel Ly-Ma, Paul Chow, Hans-Arno Jacobsen*

University of Toronto

### **10:05 Poster Session 3 and Break (San Carlos 1)**

#### **Graph Processing Applications**

Chair: Nachiket Kapre, University of Waterloo

### **Boosting the Performance of FPGA-based Graph Processor using Hybrid using Hybrid Memory Cube: A Case for Breadth First Search**

11:00

*Jialiang Zhang, Soroosh Khoram, Jing Li, UW-Madison*

### **ForeGraph: Exploring Large-scale Graph Processing on Multi-FPGA Architecture**

11:25 *Guohao Dai<sup>1</sup>, Tianhao Huang<sup>1</sup>, Yuze Chi<sup>2</sup>, Ningyi Xu<sup>3</sup>, Yu Wang<sup>1</sup>, Huazhong Yang<sup>1</sup>*

<sup>1</sup>Tsinghua University, <sup>2</sup>UCLA, <sup>3</sup>Microsoft Research Asia

### **FPGA-Accelerated Transactional Execution of Graph Workloads**

11:50 *Xiaoyu Ma<sup>1</sup>, Dan Zhang<sup>2</sup>, Derek Chiou<sup>3</sup>*

<sup>1</sup>UT-Austin, <sup>2</sup>University of Texas at Austin, <sup>3</sup>Microsoft/UT Austin

### **12:15 Lunch (Ferrantes Room, 10<sup>th</sup> Floor)**

#### **Virtualization and Applications**

Chair: John Lockwood, Algo-Logic Systems

### **Enabling Flexible Network FPGA Clusters in a Heterogenous Cloud Data Center**

2:00 *Naif Tarafdar, Thomas Lin, Eric Fukuda, Hadi Bannazadeh, Alberto Leon-Garcia,*

*Paul Chow, University of Toronto*

### **Energy Efficient Scientific Computing on FPGAs using OpenCL**

2:25 *Dennis Weller<sup>1</sup>, Fabian Oboril<sup>1</sup>, Dimitar Lukarski<sup>2</sup>, Juergen Becker<sup>1</sup>, Mehdi Tahoori<sup>1</sup>*

<sup>1</sup>Karlsruhe Institute of Technology, <sup>2</sup>PARALUTION Labs

### **Secure Function Evaluation using an FPGA Overlay Architecture**

2:50

*Xin Fang, Stratis Ioannidis, Miriam Leeser, Northeastern University*

### **3:15 Break**

## Applications

Chair: Miriam Leeser, Northeastern University

### 3:45 **FPGA Acceleration for Computational Glass-Free Displays**

*Zhuolun He and Guojie Luo, Peking University*

### 4:10 **Hardware Acceleration of the Pair-HMM Algorithm for DNA Variant Calling**

*Sitao Huang<sup>1</sup>, Gowthami Jayashri Manikandan<sup>2</sup>, Anand Ramachandran<sup>2</sup>, Kyle Rupnow<sup>3</sup>, Wen-mei W. Hwu<sup>2</sup>, Deming Chen<sup>2</sup>*

<sup>1</sup>University of Illinois at Urbana-Champaign, <sup>2</sup>UIUC, <sup>3</sup>Advanced Digital Sciences Center

### 4:35 **Conference Closing and Best Paper Award**

*Jason Anderson (University of Toronto), Jonathan Greene (Microsemi)*

## Poster Session 1

### **Measuring the Power-Constrained Performance and Energy Gap between FPGAs and Processors**

*Andy Ye<sup>1</sup> and Karthik Ganesan<sup>2</sup>*

<sup>1</sup>Ryerson University, <sup>2</sup>University of Toronto

### **A Mixed-Signal Data-Centric Reconfigurable Architecture enabled by RRAM Technology**

*Yue Zha<sup>1</sup>, Jialiang Zhang<sup>1</sup>, Zhiqiang Wei<sup>2</sup>, Jing Li<sup>1</sup>*

<sup>1</sup>UW-Madison, <sup>2</sup>Panasonic

### **A Framework for Iterative Stencil Algorithm Synthesis on FPGAs from OpenCL Programming Model**

*Shuo Wang and Yun Liang, Peking University*

### **Scala Based FPGA Design Flow**

*Yanqiang Liu<sup>1</sup>, Yao Li<sup>1</sup>, Weilun Xiong<sup>1</sup>, Meng Lai<sup>1</sup>, Cheng Chen<sup>2</sup>, Zhengwei Qi<sup>1</sup>, Haibing Guan<sup>1</sup>*

<sup>1</sup>Shanghai JiaoTong University, <sup>2</sup>Morgan Stanley

### **Thermal Flattening in 3D FPGAs using Embedded Cooling**

*Girish Deshpande and Dinesh Bhatia, UT-Dallas*

### **A Machine Learning Framework for FPGA Placement**

*Gary Grewal, Shawki Areibi, Matthew Westrik, Ziad Abuowaimer, Betty Zhao, University of Guelph*

### **Precise Coincidence Detection on FPGAs: Three Case Studies**

*Ralf Salomon and Ralf Joost, University of Rostock*

### **DTP: Enabling Exhaustive Exploration of FPGA Temporal Partitions for Streaming HPC Applications**

*Mostafa Koraei<sup>1</sup>, Magnus Jahre<sup>2</sup>, S.Omid Fatemi<sup>1</sup>*

<sup>1</sup>Tehran University, <sup>2</sup>Norwegian University of Science and Technology

**Accurate and Efficient Hyperbolic Tangent Activation Function on FPGA using the DCT Interpolation Filter**

*Ahmed Abdelsalam, Pierre Langlois, Farida Cheriet, École Polytechnique de Montréal*

**An FPGA Overlay Architecture for Cost Effective Regular Expression Search**

*Thomas Luinaud, J.M. Pierre Langlois, Yvon Savaria, École Polytechnique de Montréal*

**Storage-Efficient Batching for Minimizing Bandwidth of Fully-Connected Neural Network Layers**

*Yongming Shen, Michael Ferdman, Peter Milder, Stony Brook University*

**Poster Session 2**

**Using Vivado-HLS for Structural Design: a NoC Case Study**

*Zhipeng Zhao and James C. Hoe, CMU*

**Automatic Generation of Hardware Sandboxes for Trojan Mitigation in Systems on Chip**

*Christophe Bobda<sup>1</sup>, Taylor Whitaker<sup>1</sup>, Charles Kamhoua<sup>2</sup>, Kevin Kwiat<sup>2</sup>, Laurent Njilla<sup>2</sup>*

*<sup>1</sup>University of Arkansas, <sup>2</sup>Air Force Research Lab*

**Accelerating Financial Market Server through Hybrid List Design**

*Haohuan Fu<sup>1</sup>, Conghui He<sup>1</sup>, Huabin Ruan<sup>1</sup>, Itay Greenspon<sup>2</sup>, Wayne Luk<sup>3</sup>, Yongkang Zheng<sup>4</sup>, Junfeng Liao<sup>1</sup>, Qing Zhang<sup>4</sup>, Guangwen Yang<sup>1</sup>*

*<sup>1</sup>Tsinghua University, <sup>2</sup>Maxeler Technologies, <sup>3</sup>Imperial College London, <sup>4</sup>China Financial Futures Exchange*

**Joint Modulo Scheduling and Memory Partitioning with Multi-Bank Memory for High-Level Synthesis**

*Tianyi Lu, Shouyi Yin, Xianqing Yao, Zhicong Xie, Leibo Liu, Shaojun Wei*

*Tsinghua University*

**A Batch Normalization Free Binarized Convolutional Deep Neural Network on an FPGA**

*Hiroki Nakahara<sup>1</sup>, Haruyoshi Yonekawa<sup>1</sup>, Hisashi Iwamoto<sup>2</sup>, Masato Motomura<sup>3</sup>*

*<sup>1</sup>Tokyo Institute of Technology, <sup>2</sup>Poco a Poco Networks, <sup>3</sup>Hokkaido University*

**A 7.663-TOPS 8.2-W Energy-efficient FPGA Accelerator for Binary Convolutional Neural Networks**

*Yixing Li<sup>1</sup>, Zichuan Liu<sup>2</sup>, Kai Xu<sup>1</sup>, Fengbo Ren<sup>1</sup>, Hao Yu<sup>2</sup>*

*<sup>1</sup>Arizona State University, <sup>2</sup>Nanyang Technological University*

**CPU-FPGA Co-Optimization for Big Data Applications: A Case Study of In-Memory Samtools Sorting**

*Jason Cong<sup>1</sup>, Zhenman Fang<sup>1</sup>, Muhuan Huang<sup>2</sup>, Libo Wang<sup>1</sup>, Di Wu<sup>1</sup>*

*<sup>1</sup>UCLA, <sup>2</sup>University of California, Los Angeles*

**Stochastic-Based Multi-stage Streaming Realization of a Deep Convolutional Neural Network**

*Mingjie Lin and Mohammed Alawad, University of Central Florida*

**fpgaConvNet: Automated Mapping of Convolutional Neural Networks on FPGAs**

*Stylianos Venieris and Christos Bouganis, Imperial College London*

## Poster Session 3

### **FPGA-based Hardware Accelerator for Image Reconstruction in Magnetic Resonance Imaging**

*Emanuele Pezzotti<sup>1</sup>, Alex Iacobucci<sup>1</sup>, Gregory Nash<sup>2</sup>, Umer Cheema<sup>1</sup>, Paolo Vinella<sup>1</sup>, Rashid Ansari<sup>1</sup>*

<sup>1</sup>University of Illinois at Chicago, <sup>2</sup>University of Illinois at Chicago, Altera

### **ASAP: Accelerated Short Read Alignment on Programmable Hardware**

*Subho Banerjee, Mohamed El Hadedy, Jong Bin Lim, Daniel Chen, Zbigniew T. Kalbarczyk, Deming Chen, Ravishankar K. Iyer, UIUC*

### **RxRE: Throughput Optimization for High-Level Synthesis using Resource-Aware Regularity Extraction**

*Atieh Lotfi and Rajesh Gupta, UCSD*

### **GRT 2.0: An FPGA-based SDR Platform for Cognitive Radio Networks**

*Haoyang Wu<sup>1</sup>, Tao Wang<sup>1</sup>, Zhiwei Li<sup>1</sup>, Boyan Ding<sup>1</sup>, Xiaoguang Li<sup>1</sup>, Tianfu Jiang<sup>1</sup>, Jun Liu<sup>1</sup>, Songwu Lu<sup>2</sup>*

<sup>1</sup>Peking University, <sup>2</sup>UCLA

### **FPGA Implementation of Non-Uniform DFT for Accelerating Wireless Channel Simulations**

*Srinivas Siripurapu<sup>1</sup>, Aman Gayasen<sup>2</sup>, Nitin Chandrachoodan<sup>1</sup>, Padmini Gopalakrishnan<sup>2</sup>*

<sup>1</sup>IIT Madras, <sup>2</sup>Xilinx

### **Learning Convolutional Neural Networks for Data-Flow Graph Mapping on Spatial Programmable Architectures**

*Shouyi Yin, Dajiang Liu, Lifeng Sun, Xinhan Lin, Leibo Liu, Shaojun Wei*  
Tsinghua University

### **Cache Timing Attacks from The SoCFPGA Coherency Port**

*Sumanta Chaudhuri, Telecom ParisTech*

### **Dynamic Partitioning for Library based Placement on Heterogeneous FPGAs**

*Fubing Mao<sup>1</sup>, Wei Zhang<sup>2</sup>, Bingsheng He<sup>3</sup>, Siew Kei Lam<sup>1</sup>*

<sup>1</sup>Nanyang Technological University, <sup>2</sup>Hong Kong University of Science and Technology, <sup>3</sup>National University of Singapore

### **An Energy-Efficient Design-Time Scheduler for FPGAs Leveraging Dynamic Frequency Scaling Emulation**

*Wei Ting Loke<sup>1</sup> and Chin Yang Koay<sup>2</sup>*

<sup>1</sup>National University of Singapore, <sup>2</sup>Xilinx Asia Pacific