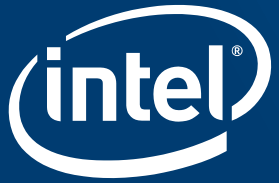

Accelerating parallel computing by closely linking of CPU and FPGA

Miroslaw Walukiewicz, Intel Corporation
September 2016

Legal Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Axxia, Intel, the Intel logo, Intel Atom, Intel Core, Intel. Experience What's Inside, the Intel. Experience What's Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Introduction

Intel's portfolio is uniquely able to fuel the agile data center in a hyper-connected world

Intel® Xeon® with integrated FPGA processors enable parallel computing and are architected to improve performance per watt

Intel® is simplifying FPGA development and delivers easy deployment inside datacenter

Key Challenges Impacting Parallel Computing in Data Centers

Increasing Velocity of Unique Workloads
Homogeneity vs. Customization

Power Consumption
Densely packed processing and acceleration

One Server – Multiple Use-cases
Reuse same hardware for various purposes

What is an FPGA?

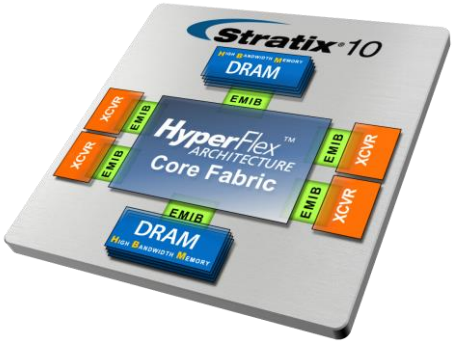

- A user-customizable semiconductor
- Functions can be reprogrammed as market dynamics change or standards evolve
- Can improve performance per watt, bandwidth & latency
- FPGAs can enhance CPU-based processing by *accelerating algorithms* and *minimizing bottlenecks for a wide variety of workloads*

FPGA = Field-programmable gate array
XCVR = High Speed Serial Transceiver Block

Intel® Xeon® with integrated FPGA Overview

Discrete FPGA

Discrete FPGA with Xeon®

Intel® Xeon® with integrated FPGA Platform

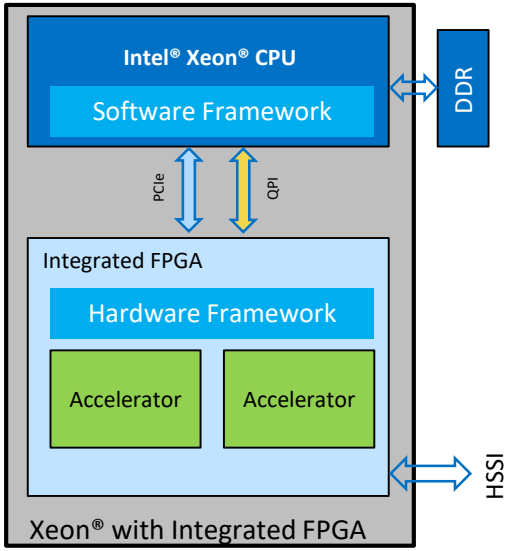

BDX+FPGA MCP Pilot

HDL Programming

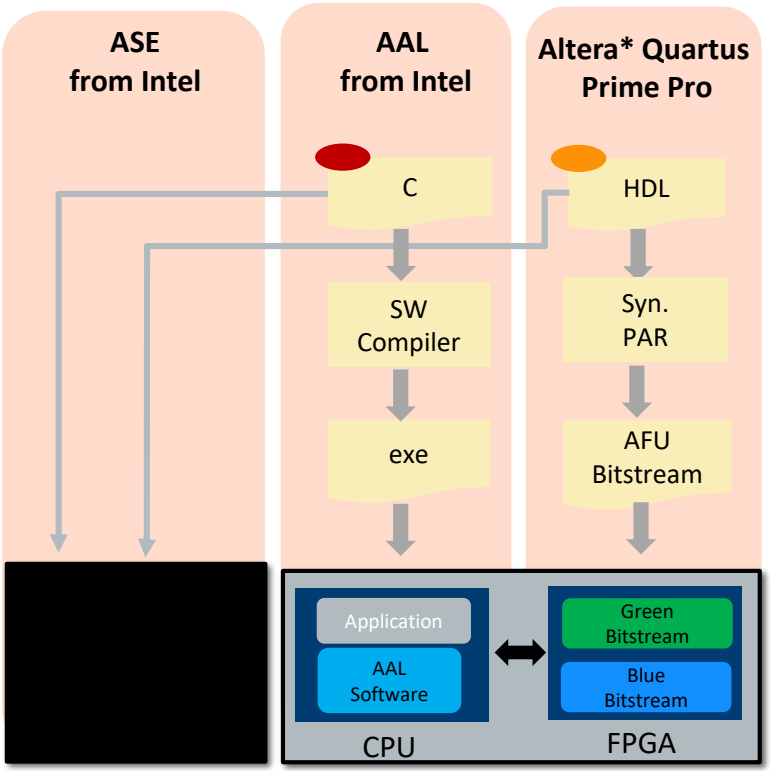

OpenCL* Programming

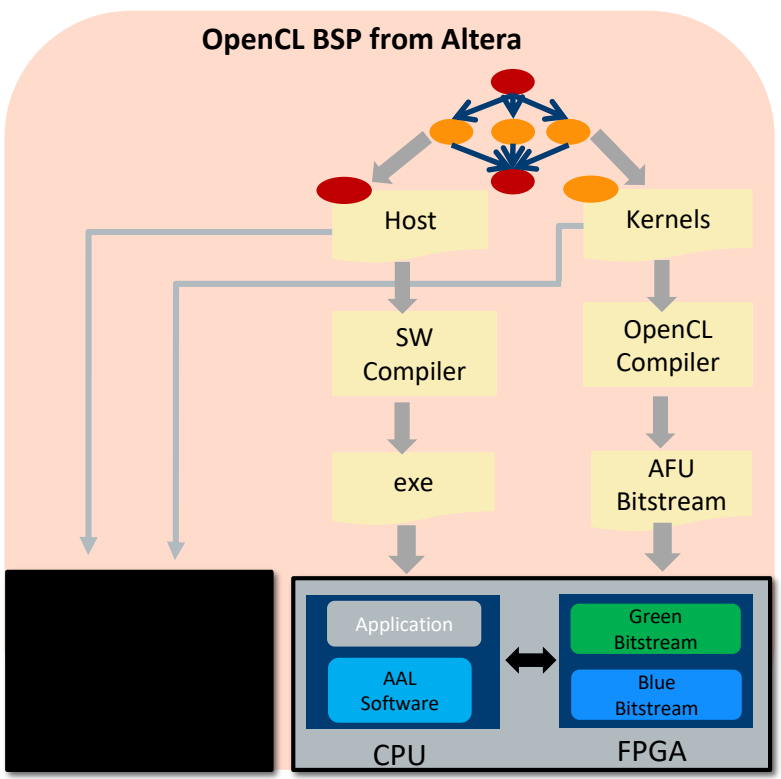

Intel® Xeon+FPGA Software Stack

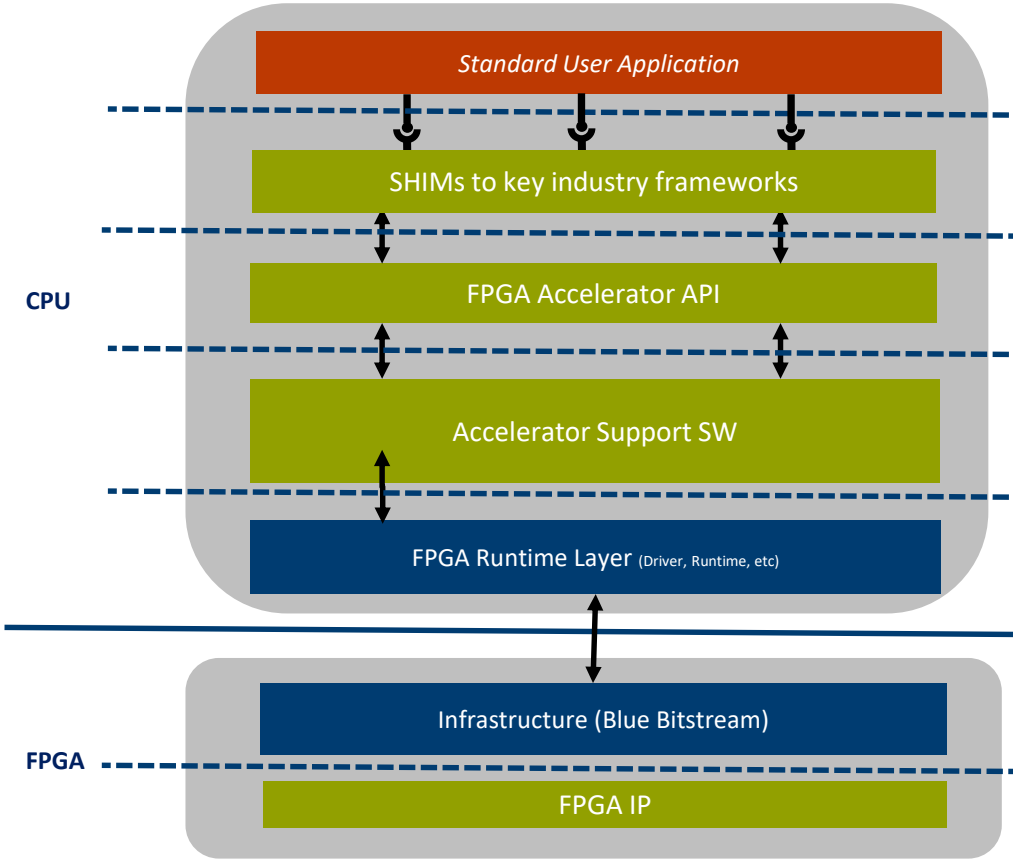

User Application
* Connect to FPGA accelerator using standard interface just by pointing to FPGA libraries.

SHIM
Features
* Wrapper to connect FPGA accelerators to standard libraries

Accelerator API
Features
* Low level FPGA specific APIs
* Load balancer to efficient parallelize between FPGA and CPU to achieve max performance and accelerator utilization

Accelerator Support SW
Features
* Accelerator specific Pre/Post routine.
* CPU optimized SW implementations.
* Accelerator specific initialization, error and buffer management routines

Intel SW Infrastructure
Features
* FPGA Runtime
* Kernel drivers
* Resource Manager

Intel HW Infrastructure
Features
* Virtualization
* Remote Debug
* Security
* Temp/Pwr Monitoring
* Partial Reconfiguration,

Accelerator implementation in RTL
(example CNN, GEMM, etc.)

Intel® Xeon+FPGA Software Stack

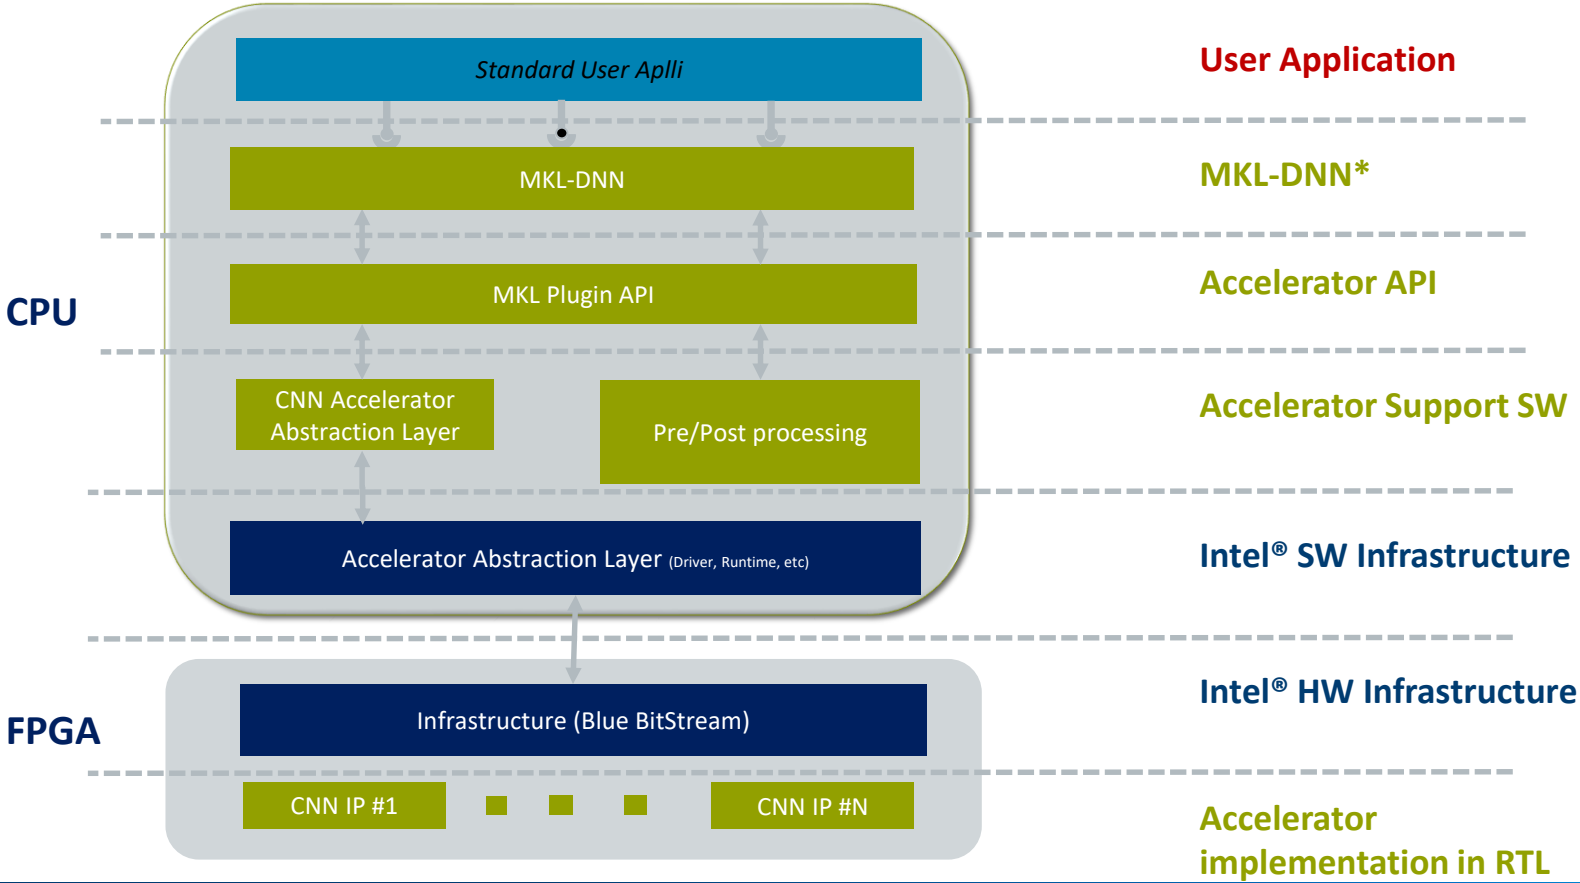

Intel® Xeon® with integrated FPGA Pilot System Examples

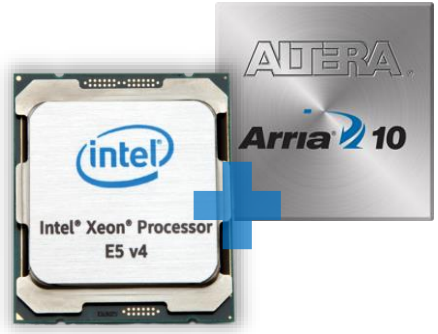

Delivering Custom acceleration

Applications	Image Identification	Security	Firewall, VPN, Router...	Xeon + FPGA
				
Algorithms	Convolutional Neural Network	Encryption	Virtual Switching	

Stochastic Gradient Descent computation

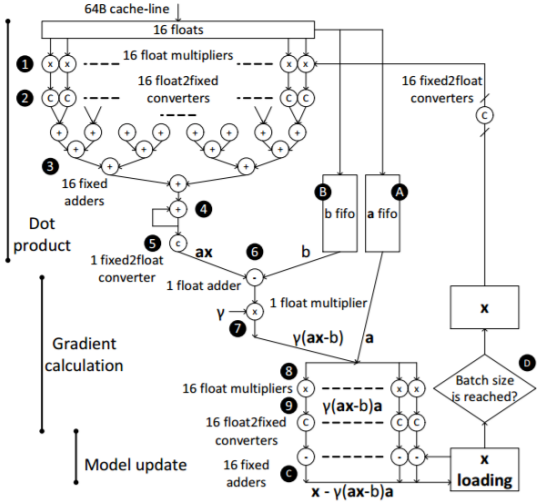

Fig. 2: Computation pipeline for `floatFSGD`, latency: 36 cycles, data width: 64B, processing rate: 64B/cycle.

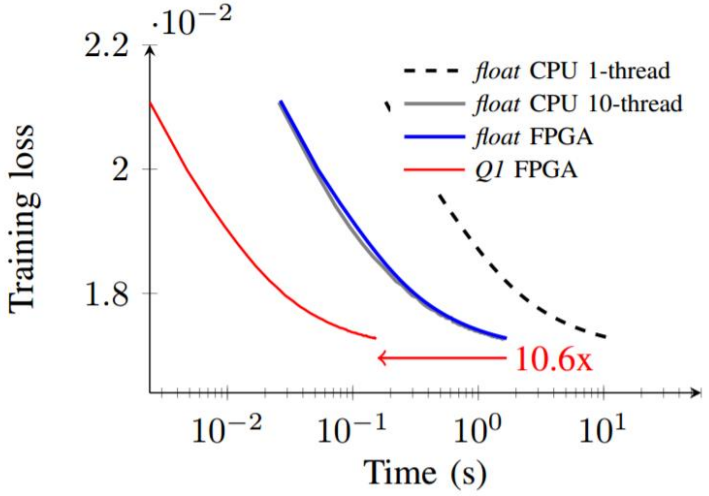

(c) *mnist*, trained for digit 7. $\gamma = 1/2^{15}$.

FPGA-accelerated Dense Linear Machine Learning: A Precision-Convergence Trade-off Kaan Kara∗ , Dan Alistarh∗†, Gustavo Alonso∗ , Onur Mutlu∗ , Ce Zhang∗ ∗ Systems Group, Department of Computer Science ETH Zurich, Switzerland firstname.lastname@inf.ethz.ch † Distributed Algorithms and Systems IST, Austria

Hardware Accelerator Research Program

- Global program enabling access to Intel® Xeon® with integrated FPGA platforms
- More than 150 universities (most outside US) over the world involved
- 2 hosting sites with 10 platforms available
- Most of the participants works remotely
- Intel® hosts trainings to accepted program participants

Global community, across continents

Hardware Accelerator Research Program – areas of interest

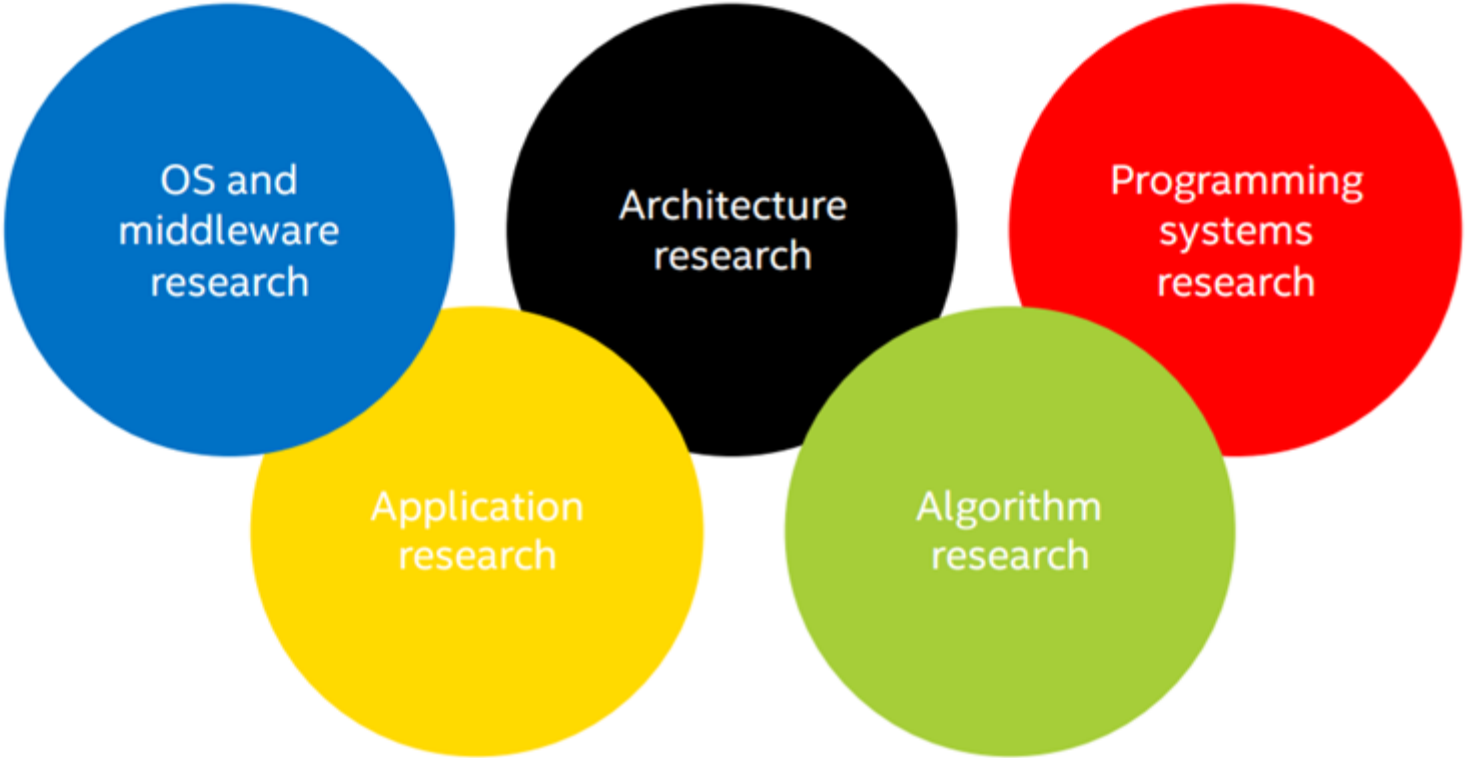

Send us your proposals by hw_accelerator_research_program@intel.com

You are invited to join

Hardware Accelerator Research Program

Thank you!

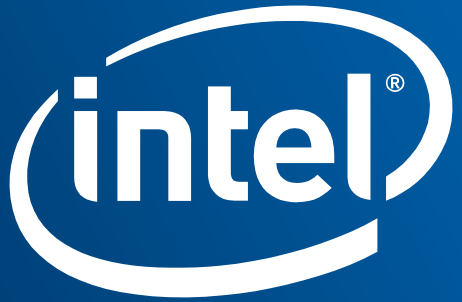
intel