Intel® Xeon®+FPGA Platform for the Data Center

FPL'15 Workshop on Reconfigurable Computing for the Masses

PK Gupta, Director of Cloud Platform Technology, DCG/CPG

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- Applications and Eco-system

Digital Services Economy…

50[1] Billion **DEVICES**

New **SERVICES** $450B[2]

Build out of the **CLOUD** $120B[3]

1: Sources: AMS Research, Gartner, IDC, McKinsey Global Institute, and various others industry analysts and commentators
2: Source IDC, 2013. 2016 calculated base don reported CAGR '13-'17
3: Source: IDATA /Digiworld, 2013

...Fueling Cloud Computing Growth

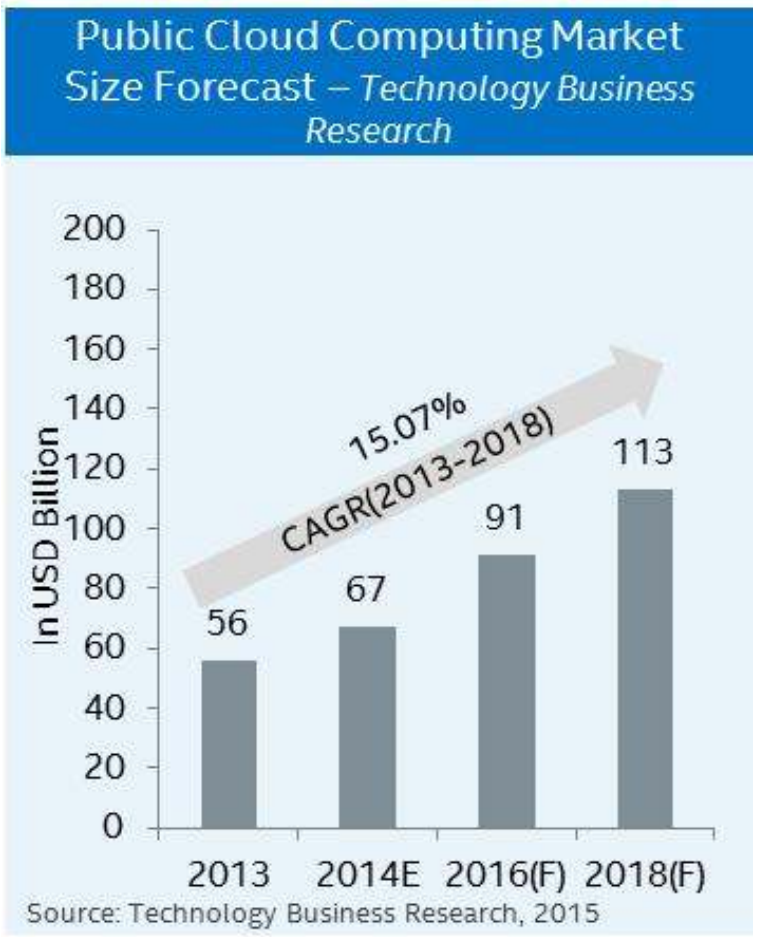

Cloud Economics

Workload Performance Metrics

- VMs per System
- Web Transactions / Sec
- Storage Capacity
- Hadoop Queries

÷

Amazon's TCO Analysis[1]

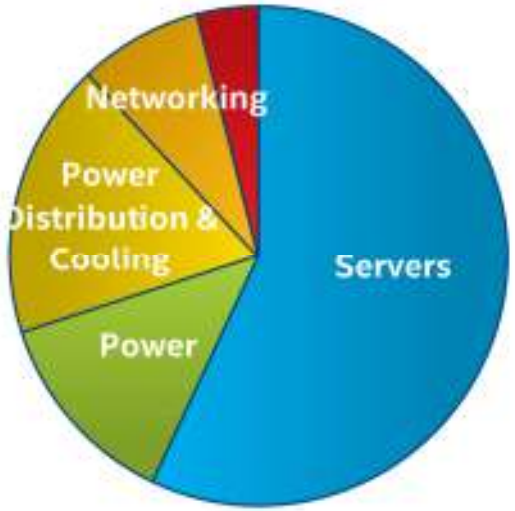

Performance / TCO is the key metric

1: Source: James Hamilton, Amazon* http://perspectives.mvdirona.com/2010/09/overall-data-center-costs/

Diverse Data Center Demands

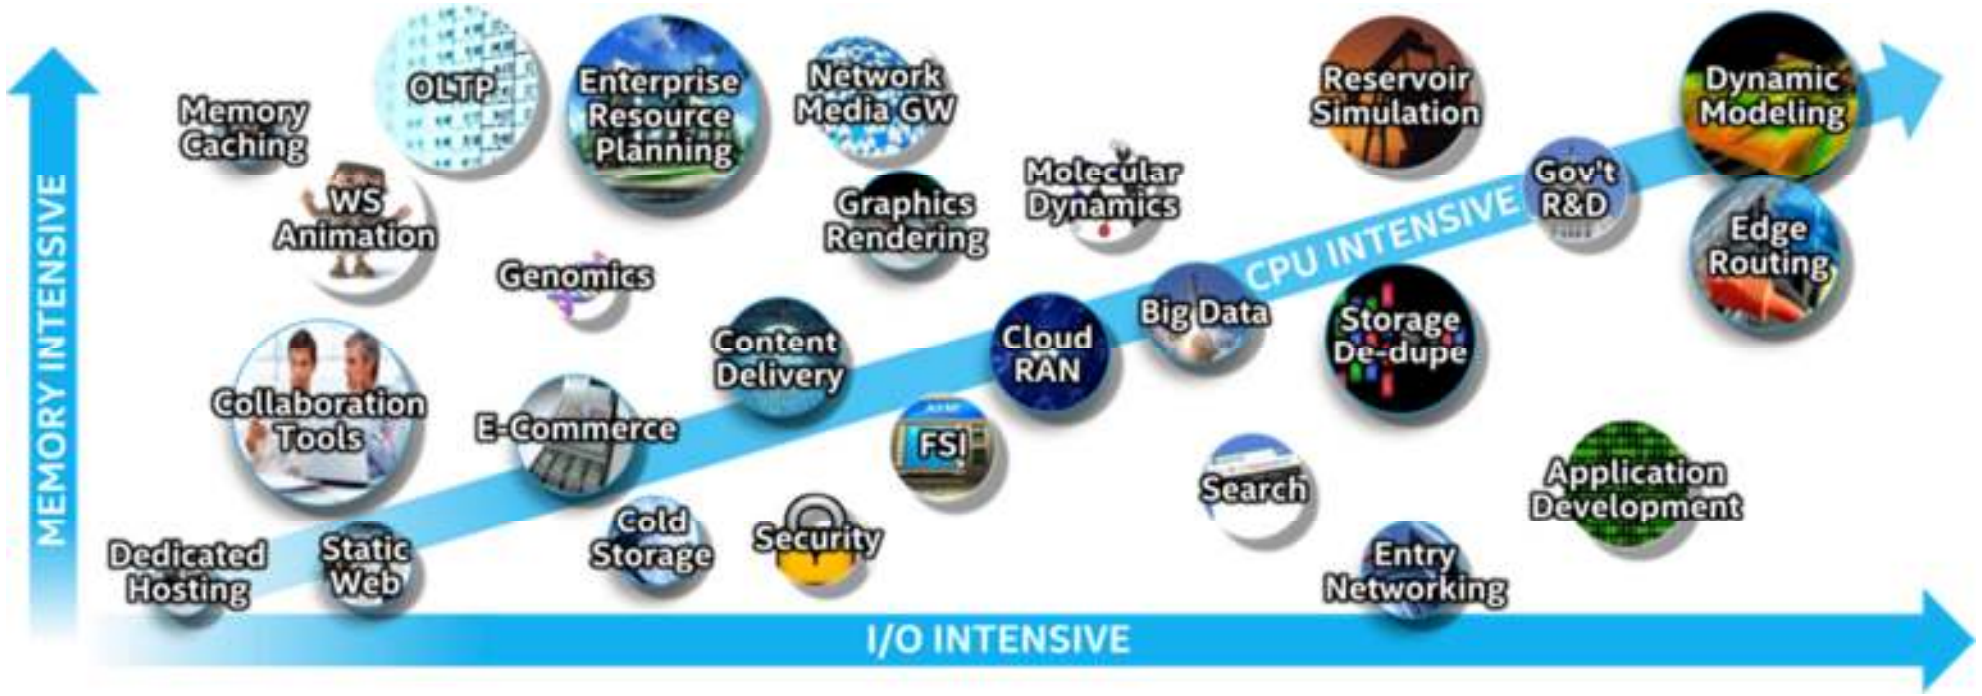

Accelerators can increase Performance at lower TCO for targeted workloads

Intel estimates; bubble size is relative CPU intensity

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- Applications and Eco-system

Accelerator Architecture Landscape

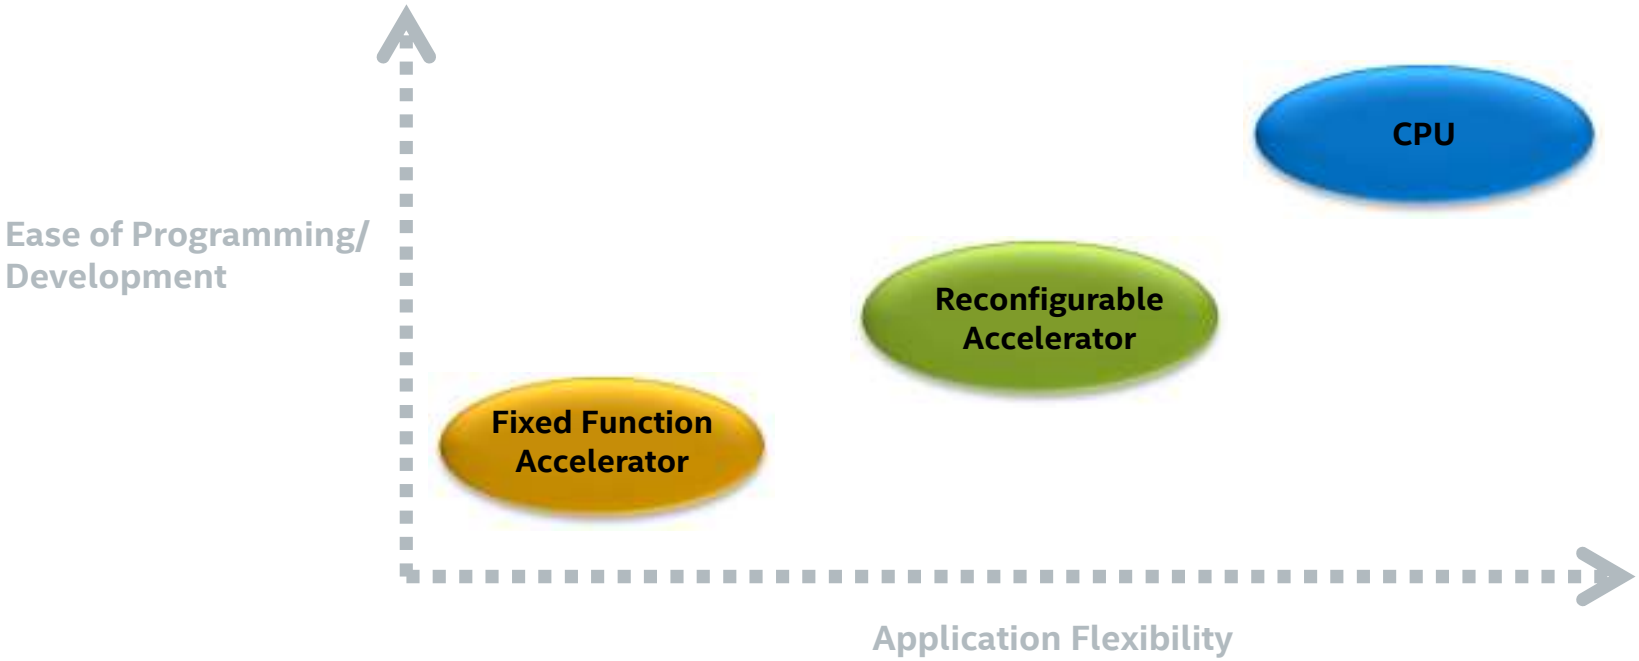

Accelerator Attach

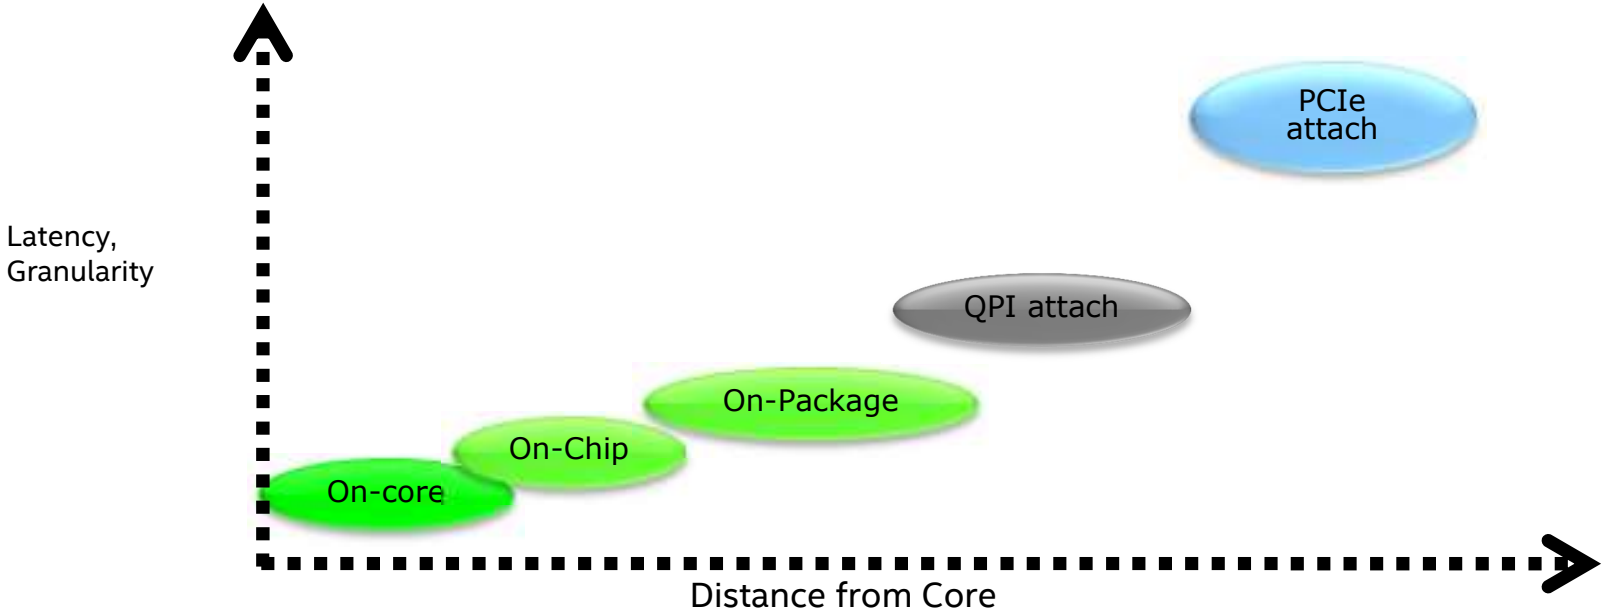

Best attach technology might be application or even algorithm dependent

Coherency and Programming Model

- Data Movement
  - In-line
    - Accelerator processes data fully or partially from direct I/O
  - Shared Virtual Memory :
    - Virtual addressing eliminates need for pinning memory buffers
    - Zero-copy data buffers
- Interaction between Core and Accelerator
  - Off-load
  - Hybrid : algorithm implemented on host and accelerator

Proposed Platform for the Data Center

- FPGA with coherent low-latency interconnect:
  - Simplified programming model
    - Support for virtual addressing
    - Data Caching
  - Enables new classes of algorithms for acceleration with:
    - Full access to system memory
    - Support for efficient irregular data pattern access
  - Remapping of algorithms from off-load model to hybrid processing model
    - Fine grained interactions

IVB+FPGA Software Development Platform

Software Development for Accelerating Workloads using Xeon and coherently attached FPGA in-socket

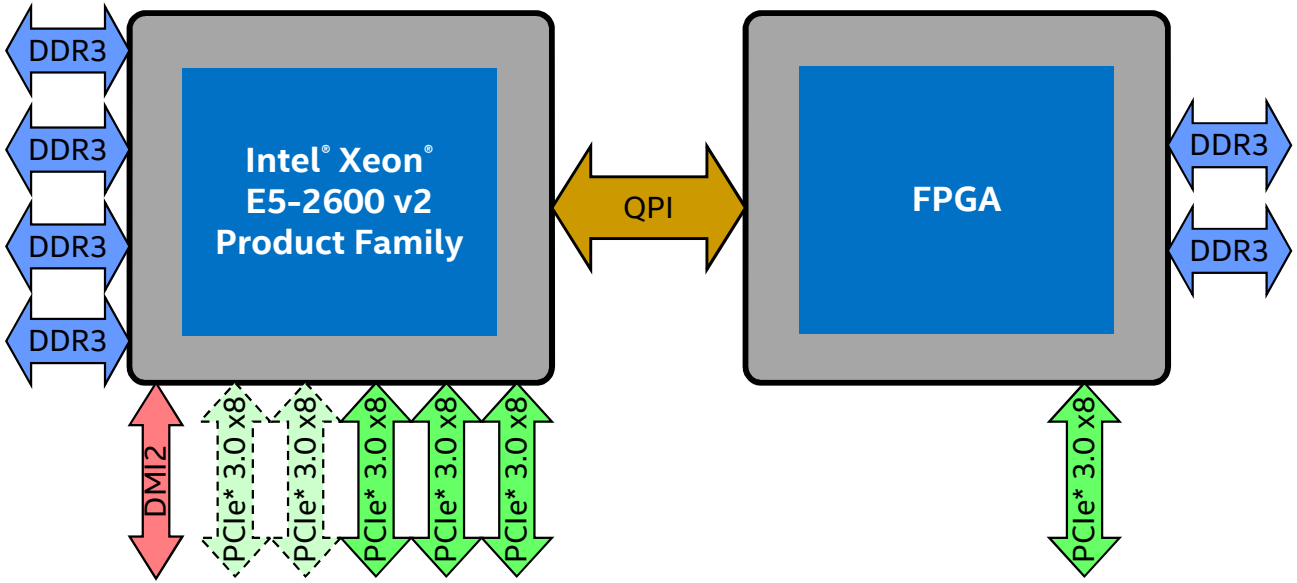

Processor	Intel® Xeon® E5-26xx v2 Processor
FPGA Module	Altera Stratix V
QPI Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR3 (up to 64 GB)
Expansion connector to FPGA Module	PCIe 3.0 x8 lanes - maybe used for direct I/O e.g. Ethernet
Features	Configuration Agent, Caching Agent,, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

Heterogeneous architecture with homogenous platform support

System Logical View

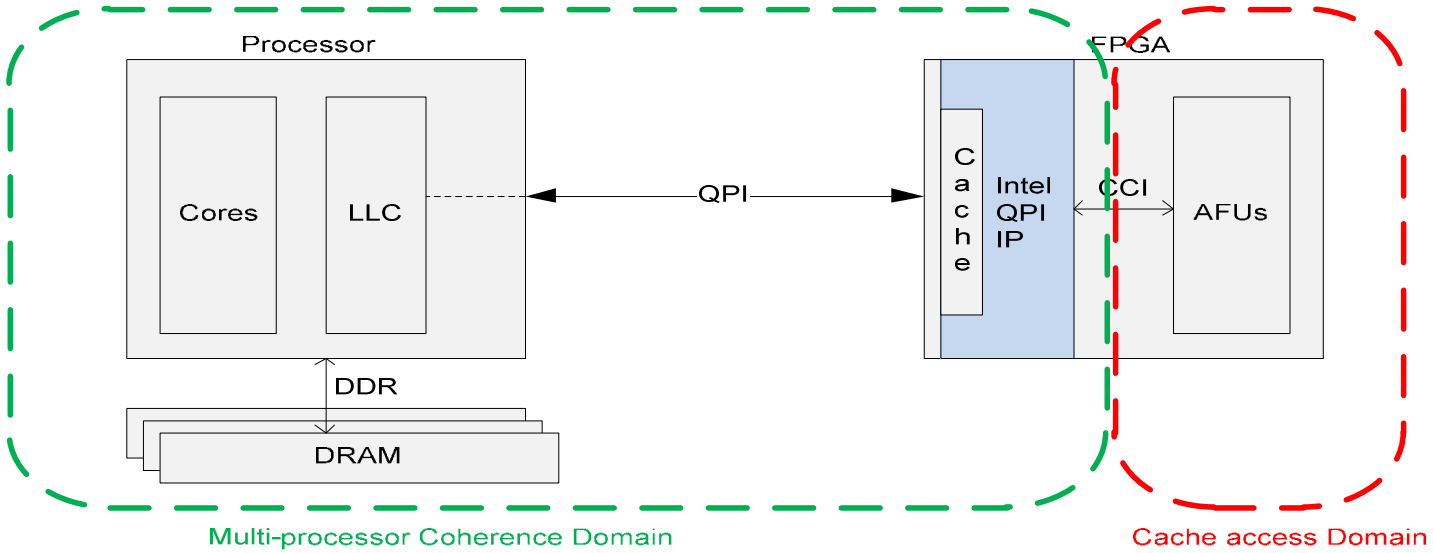

AFUs can access coherent cache on FPGA

AFUs can “not” implement a second level cache

Intel® Quick Path Interconnect (Intel® QPI) IP participates in cache coherency with Processors

Programming Interfaces

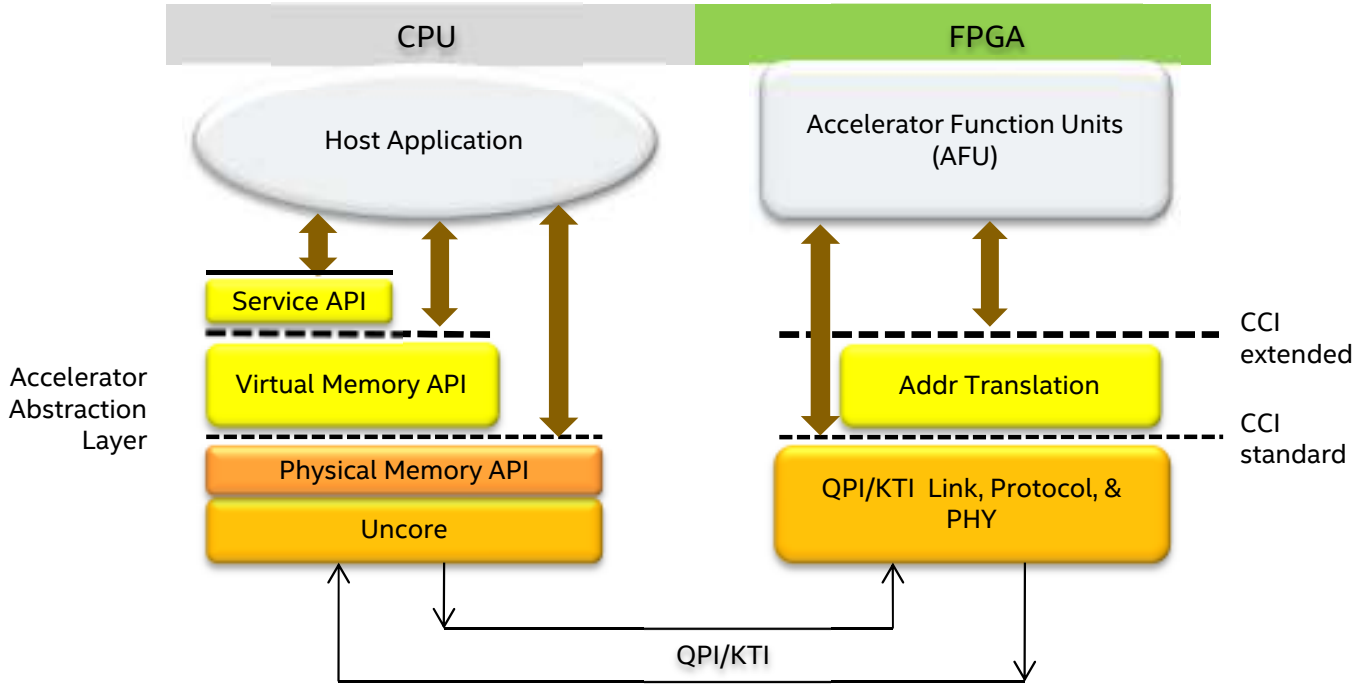

Programming interfaces will be forward compatible from SDP to future MCP solutions
Simulation Environment available for development of SW and RTL

Programming Interfaces : OpenCL

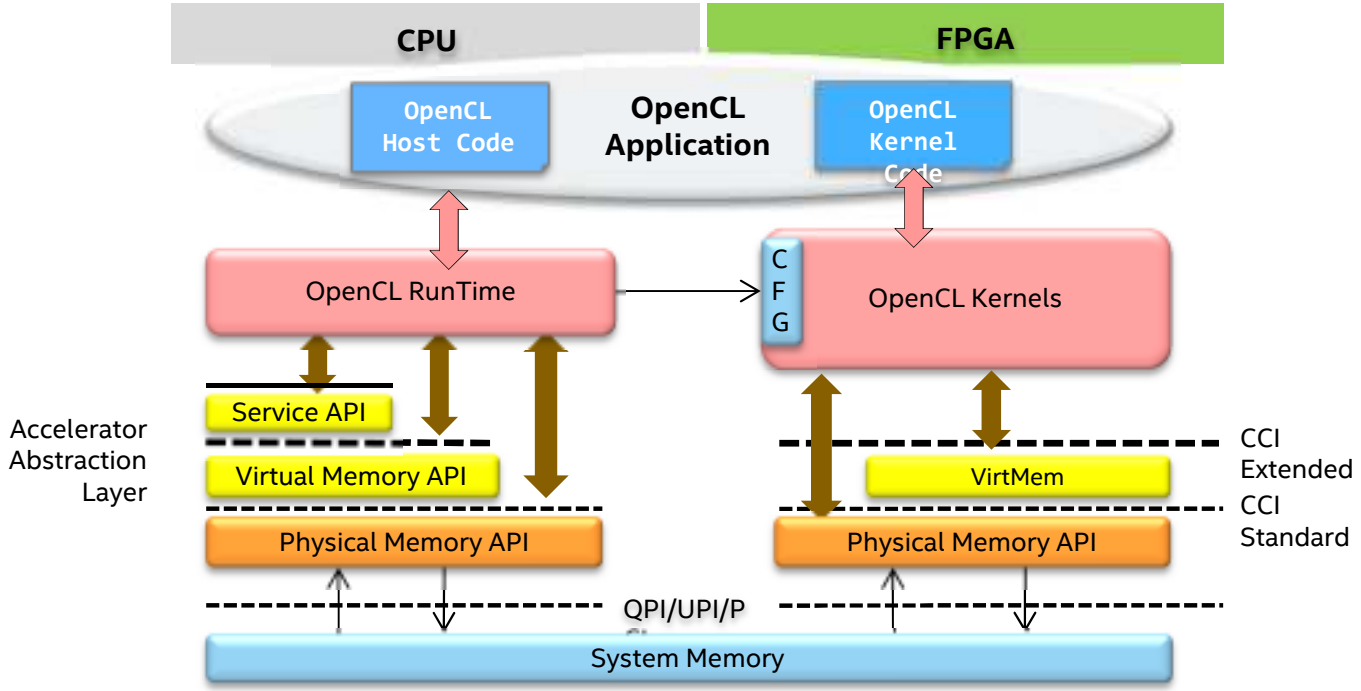

Unified application code abstracted from the hardware environment
Portable across generations and families of CPUs and FPGAs

Overview

- Data Center and Workloads
- Xeon+FPGA Accelerator Platform
- Applications and Eco-system

Intel® Xeon® + FPGA[1] in the Cloud Vision

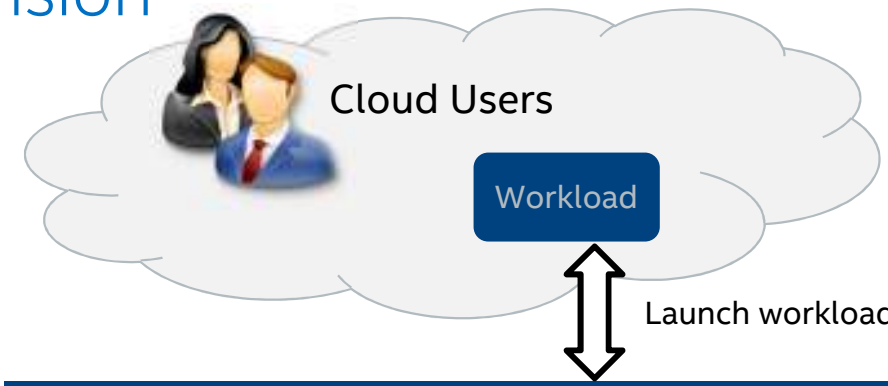

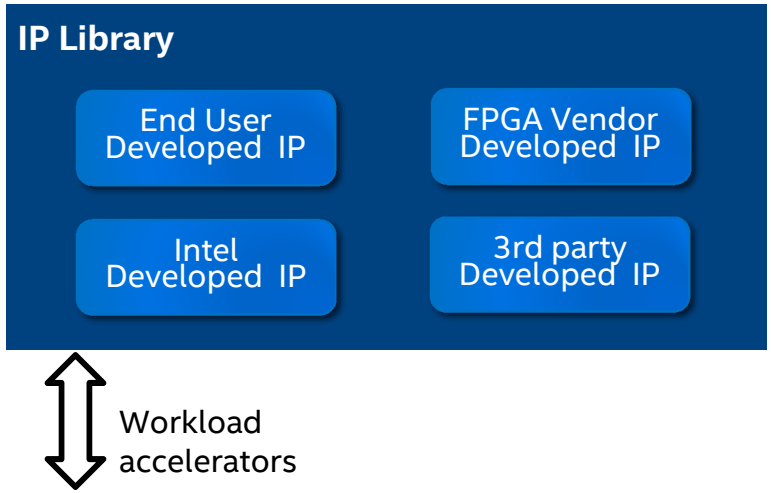

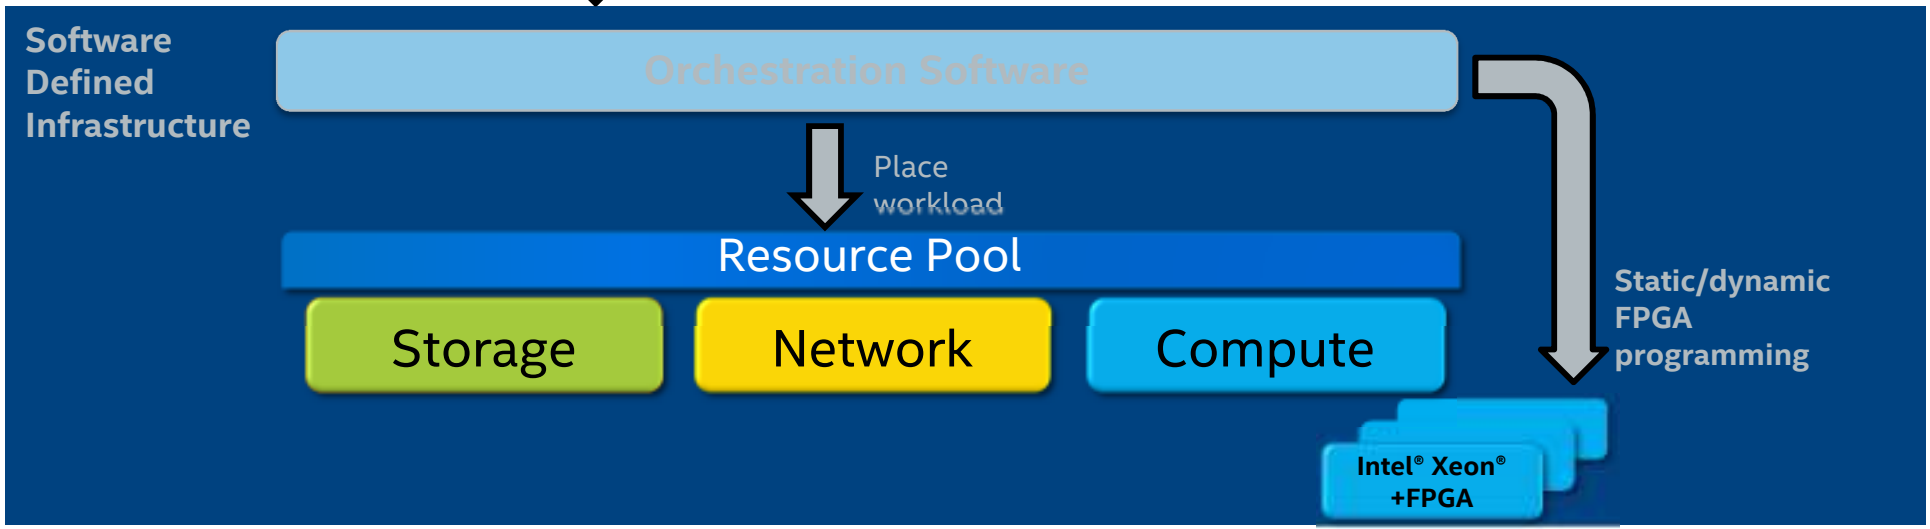

1: Field Programmable Gate Array (FPGA)

Example Usage:

Deep Learning Framework for Visual Understanding

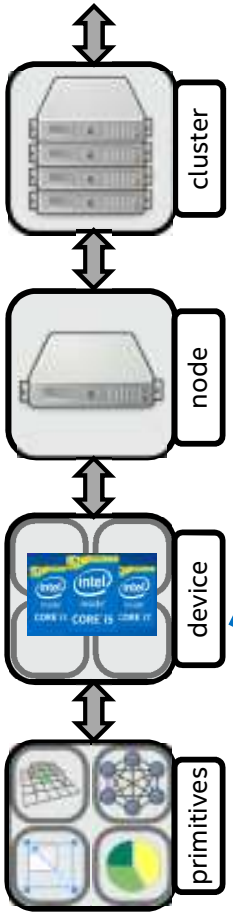

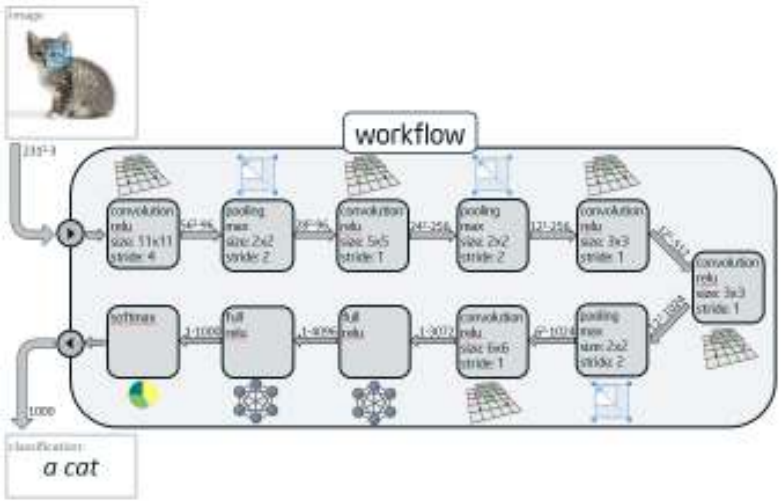

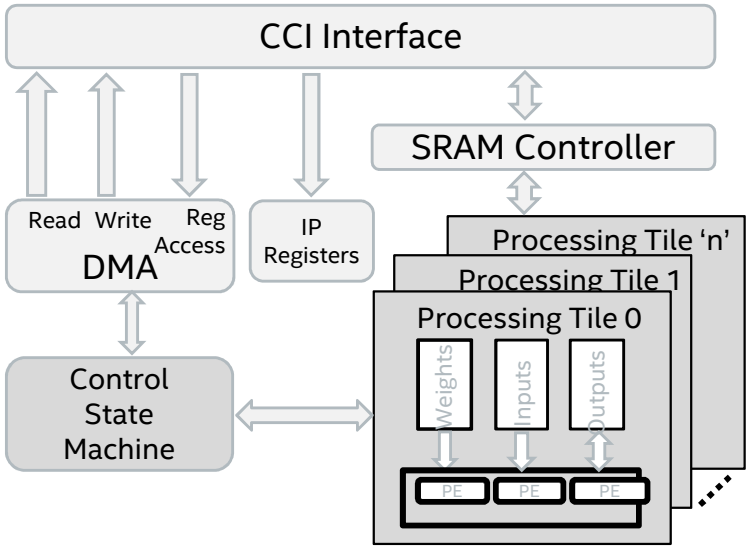

CNN (Convolutional Neural Network) function accelerated on FPGA:
*Power-performance of CNN classification boosted up to **2.2X**†*

†Source: Intel Measured (Intel® Xeon® processor E5-2699v3 results; Altera Estimated (4x Arria-10 results)
2S Intel(Xeon E5-2699v3 + 4x GX1150 PCI Express® cards. Most computations executed on Arria-10 FPGA's, 2S Intel Xeon E5-2699v3 host assumed to be near idle, doing misc. networking/housekeeping functions. Arria-10 results estimated by Altera with Altera custom classification network. 2x Intel Xeon E5-2699v3 power estimated @ 139W while doing "housekeeping" for GX1150 cards based on Intel measured microbenchmark. In order to sustain ~2400 img/s we need a I/O bandwidth of ~500 MB/s, which can be supported by a 10GigE link and software stack

Example Usage:
Genomics Analysis Toolkit

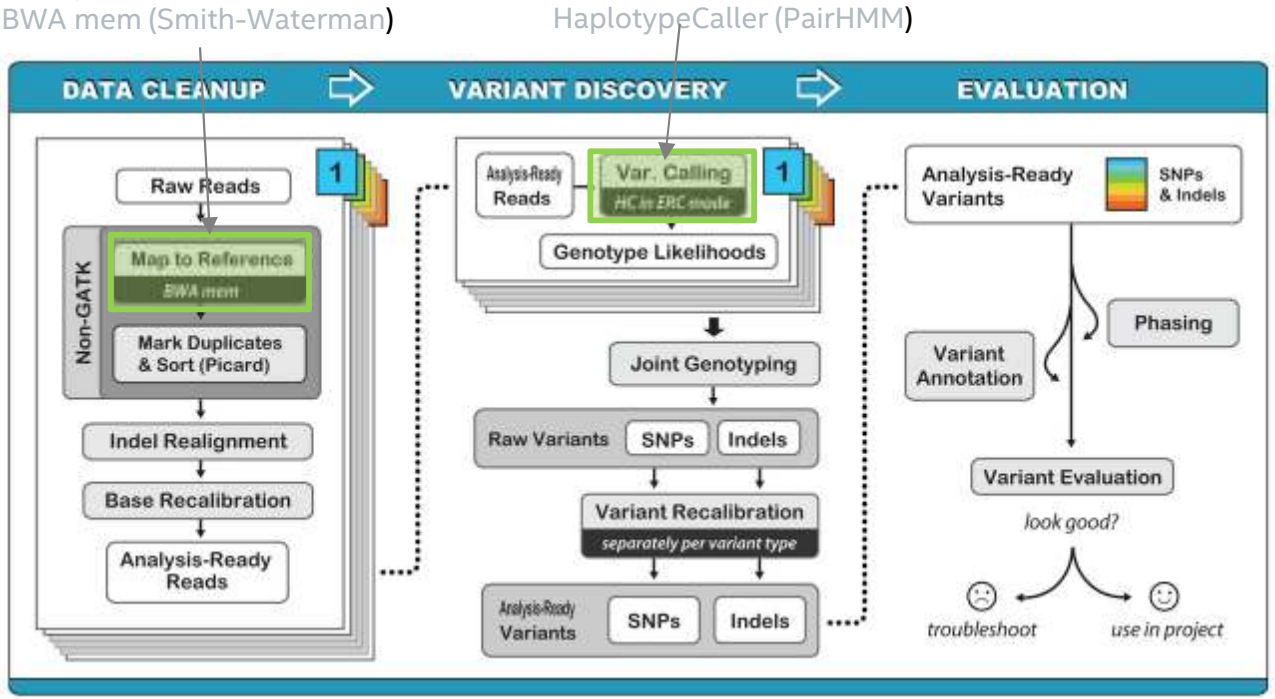

PairHMM function accelerated on FPGA:
*Power-performance of pHMM boosted up to **3.8X**†*

†pHMM Algorithm performance is measured in terms of Millions Cell Updates per seconds (CUPS).
Performance projections: **CPU Performance: includes:** 1 core Intel® Xeon® processor E5-2680v2 @ 2.8GHz delivers 2101.1 MCUP/s measured; estimated value assumes linear scaling to 10 Cores on Xeon ES2680v2 @ 2.8 GHz & 115W TDP; **FPGA Performance includes**: 1 FPGA PE (Processing Engine) delivers 408.9 MCUP/s @ 200 MHz measured; estimated value assumes linear scaling to 32 PEs and 90% frequency scaling on Stratix-V A7 400 MHz based on RTL Synthesis results (35W TDP). Intel estimated based on 1S Xeon E5-2680v2 + 1 Stratix-V A7 with QPI 1.1 @ 6.4 GT/s full width using Intel® QuickAssist FPGA System Release 3.3, ICC (CPU is essentially idle when work load is offloaded to the FPGA)

Example Usage:
Database Query Processing

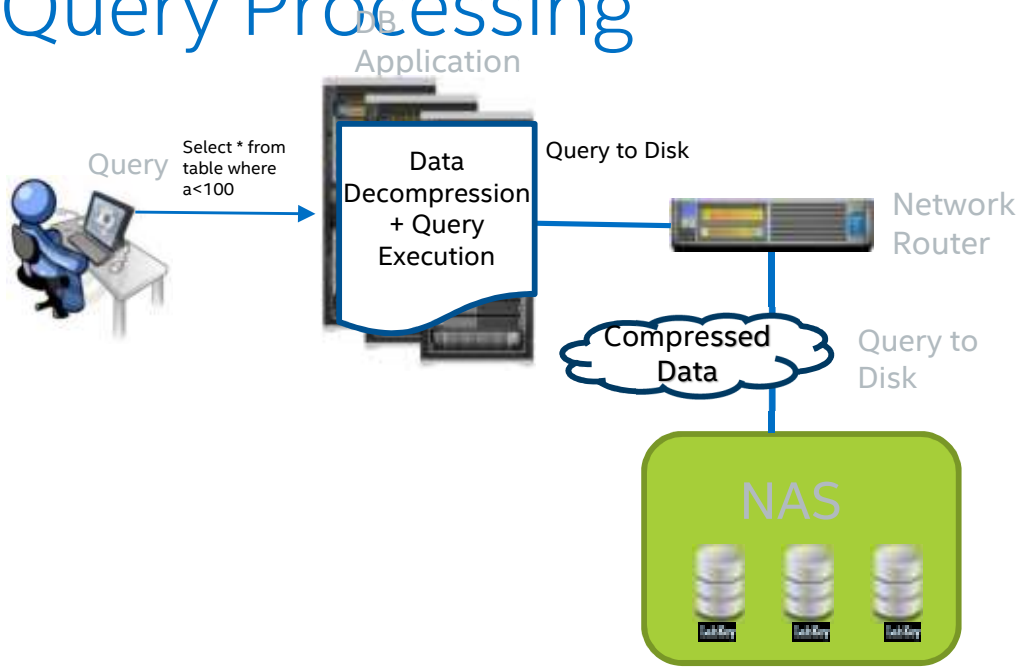

Decompression function accelerated on FPGA:
*Power-performance of LZO Decompression boosted up to **1.9X**†*

†LZO Decompression performance is measure in terms of Byte Decompressed per second.
Performance projections for stream files of size 111kB where the decompression matches are in range of FPGA buffer not requiring any system memory R/W requests: **FPGA performance (estimated):** 0.48 Clocks/Byte per LZOD PE (Processing Engine) (resulting in 727 MB/s throughput @ 350 MHz) based on cycle accurate RTL simulation measurements; assuming linear scaling to 20 LZOD PE on Arria-10 1150 @ 350 MHz (60W TDP) (CPU is essentially idle when work load is offloaded to the FPGA). **CPU performa**nce: 4.5 Clocks/Byte measured on one thread E5-2699v3 using IPP 9.0.0 (resulting in 511 MB/s Throughput @ 2.3GHz); assuming linear scaling to 36 Threads on 1S E5-2699v3 @ 2.3 GHz (145W TDP)

Academic Research in FPGA Usages

Intel & Altera jointly launched Hardware Accelerator Research Program

Q1'15: Call for proposals "which will provide faculty with computer systems containing Intel microprocessors and an Altera* Stratix* V FPGA module that incorporates Intel® QuickAssist Technology in order to spur research in programming tools, operating systems, and innovative applications for accelerator-based computing systems"

Q2'15: Proposals reviewed and selected

Q3'15: Systems being shipped to universities

Q & A