

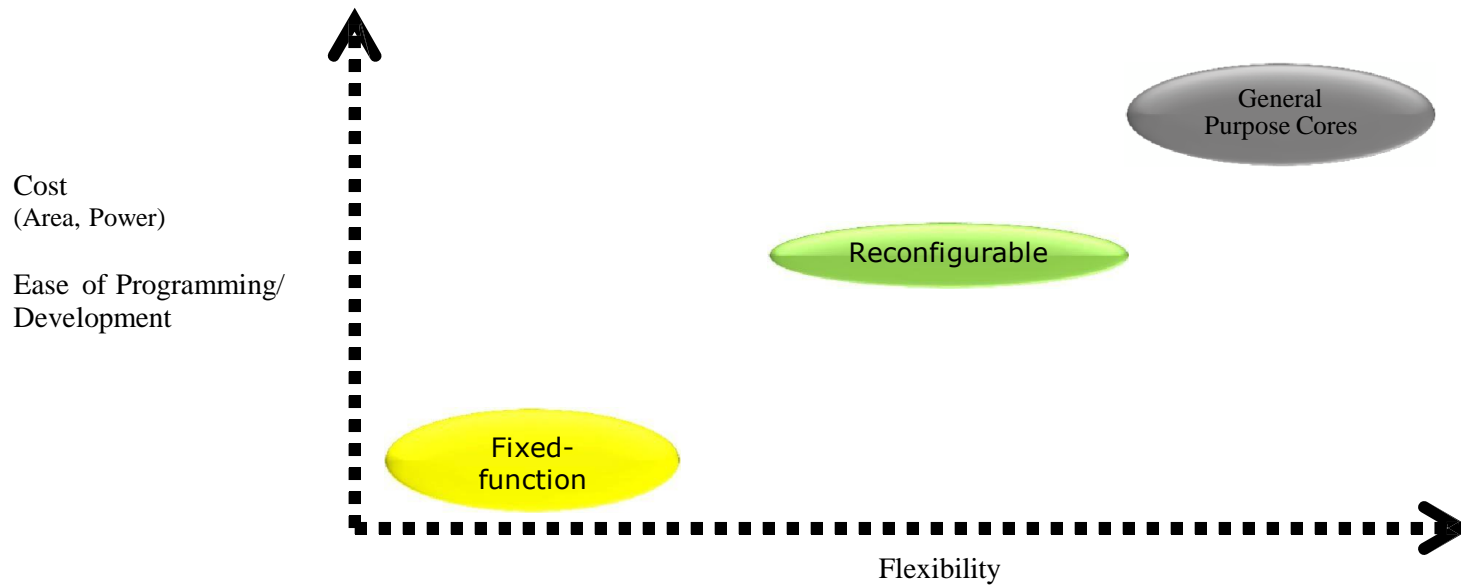
Pactron FPGA Accelerated Computing Solutions

Intel[®] Xeon + Altera FPGA

Motivation for Accelerators

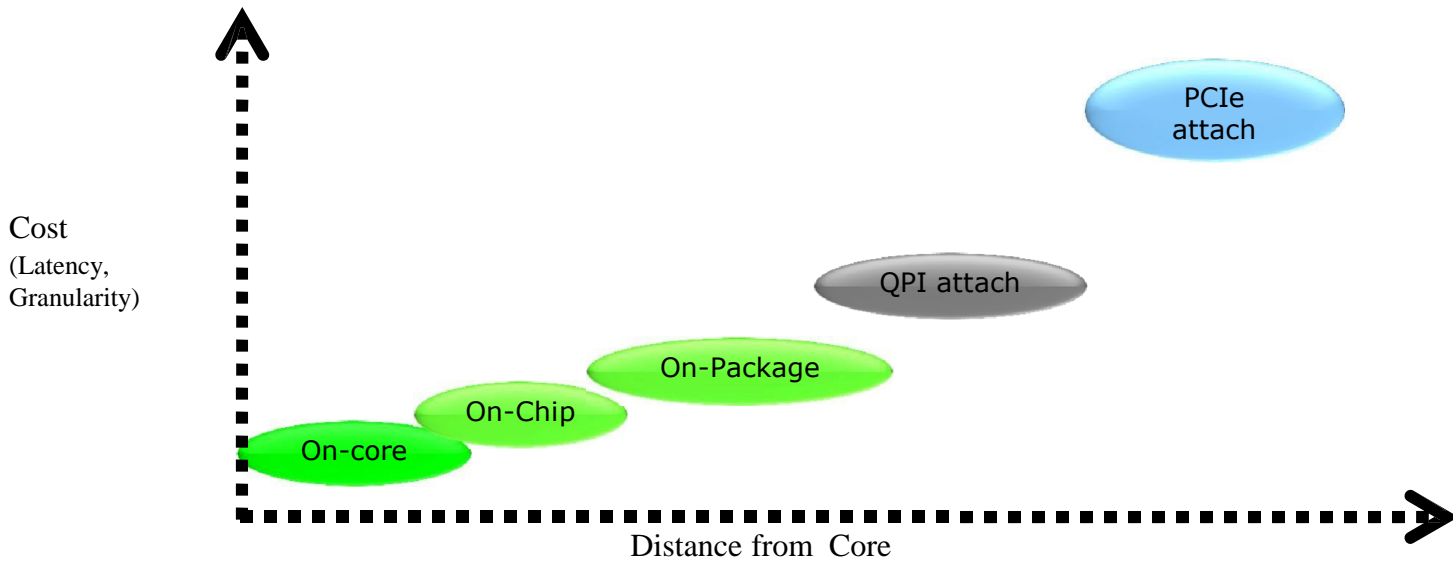
- Enhanced Performance: Accelerators compliment CPU cores to meet market needs for performance of diverse workloads in the Data Center:
 - Enhance single thread performance with tightly coupled accelerators or compliment multi-core performance with loosely coupled accelerators via PCIe or QPI attach
- Move to Heterogeneous Computing: Moore's Law continues but demands radical changes in architecture and software.
 - Architectures will go beyond homogeneous parallelism, embrace heterogeneity, and exploit the bounty of transistors to incorporate application-customized hardware.

Accelerator Architecture



Performance Efficiency: Performance/Watt, Performance/\$
Programming Complexity : Effort, Cost

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



Pactron FPGA Accelerated Computing Solutions


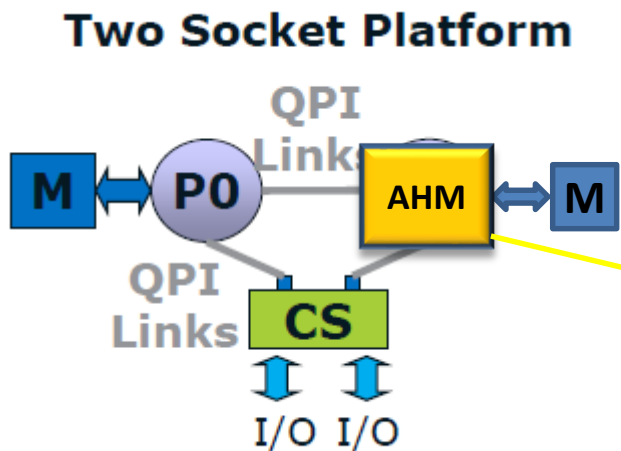
“Intel® Xeon + Altera FPGA”
Software Development Platforms

Pactron's Intel® Xeon + Altera FPGA SDP Platforms


- **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

Pactron's Intel® Xeon + Altera FPGA SDP Platforms

QPI Attached Accelerator Hardware Module ~ AHM



Intel® Xeon CPU e-2670 v2



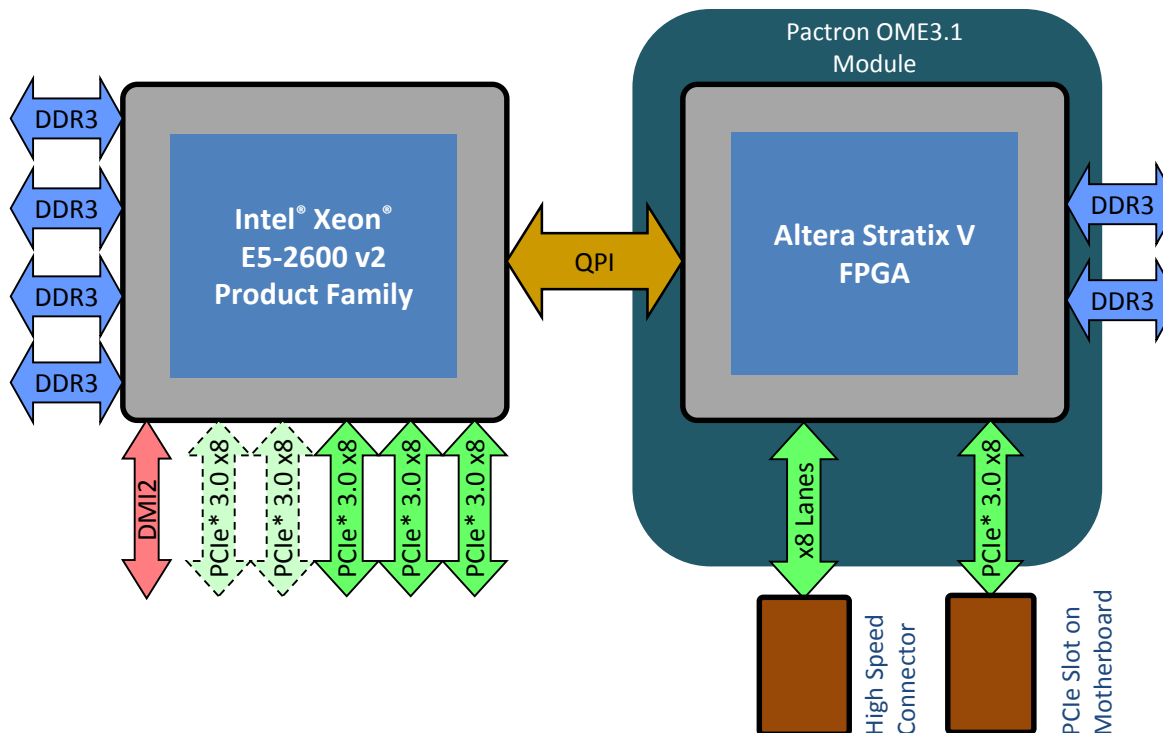
AHM



Pactron Alter FPGA Modules

Pactron's Romley "IVY Bridge" SDP Platforms

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

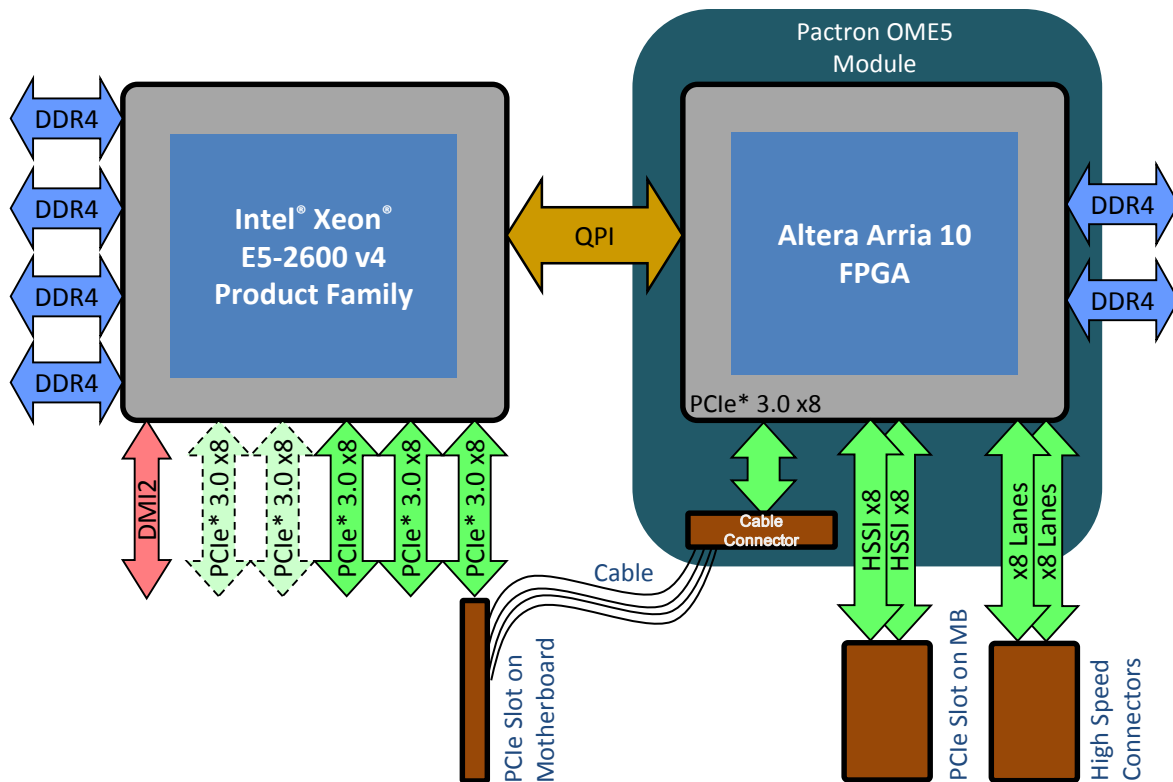


Processor	Intel® Xeon® E5-26xx v2 Processor
FPGA Module	Pactron OME3.1
QPI Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR3 (up to 64 GB)
Expansion connector to FPGA Module	PCIe 3.0 x8 lanes - maybe used for direct I/O e.g. Ethernet
Features	Configuration Agent, Caching Agent, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

Released and Shipping today

Pactron's Grantley "HSX/BSX" SDP Platforms

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



Processor	Intel® Xeon® E5-26xx v4 Processor
FPGA Module	Pactron OME5
QPI Speed	6.4 GT/s full width (target 8.0 GT/s at full width)
Memory to FPGA Module	2 channels of DDR4 (up to 64 GB)
Expansion connector to FPGA Module	Two x8 HSSI lanes (PCIe electrical) - maybe used for direct I/O e.g. Ethernet
Control Port	PCIe 3.0 x8 connected to cable connector on module. For discovery/enumeration only, not for bulk data.
Features	Configuration Agent, Caching Agent,, (optional) Memory Controller
Software	Accelerator Abstraction Layer (AAL) runtime, drivers, sample applications

**** Available Dec 2015 ****

Pactron's

Romley vs. Grantley Hardware differences

Description	Romley - OME 3.1	Grantley - OME 5.0
CPU	IVT	HSX, BDX
Platform Name	Grizzly Pass	Wild Cat Pass
Coherent Interconnect	One QPI @ 6.4GTs	One QPI @ 6.4GTs (Target 8.0GTs)
DDR Interface	2 channels of 72-bit wide DDR3. Each channel with 2 Dual Rank RDIMMs, 16GB each. DDR3 clock speed == 400 MHz	2 channels of 72-bit wide DDR4. Each channel with 2 Dual Rank RDIMMs, 16GB each. DDR4 clock speed == 800 MHz
PCIe connections to FPGA	N/A	One PCIe x8 End Point on the left edge of the FPGA along with QPI
HSSI ports	One PCIe x8 Gen 3 Root Complex on the right edge going to a PCIe slot on the platform	Two PCIe x8 Gen 3 Root Complex ports on the right edge going to PCIe slots on the platform
HSSI connectors	One High Speed connector on the OME3 module with 8 transceivers	Two High speed connectors on the OME5 module with 8 transceivers each.

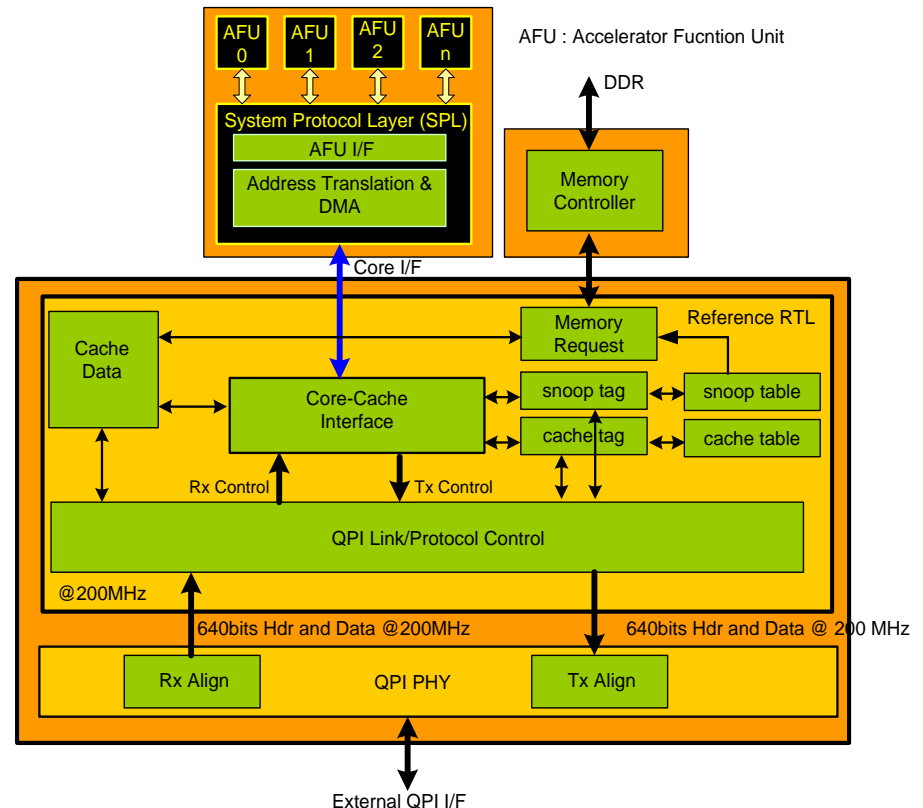
Intel QPI Reference RTL

- **PHY** – Implements the QPI PHY 1.1 (Analog/Digital)
- **QPI Link/Protocol** – Combined unit that implements QPI Link/Protocol functionality for **Cache Agent + Home Agent**
- **Core-Cache Interface (CCI)** – Implements the Core request functionality, manages Cache/Tag/Coherency interactions.
- **Cache Data** – Holds cache data
- **Cache Tag** – Tracks state of cached cacheline (MESI + internal states)
- **Snoop Tag** – Tracks state of cacheline fetched by other agents
- **Coherency/Snoop Table** – Programmable table that allows for easy modification of coherency protocol/rules
- **System Protocol Layer** – Implements DMA/Address translation functionality for Accelerator developers (Not part of Reference RTL code)
- **AFU** – Accelerator Function Unit implements acceleration logic. For Accelerator developers only. (Not part of Reference RTL code)

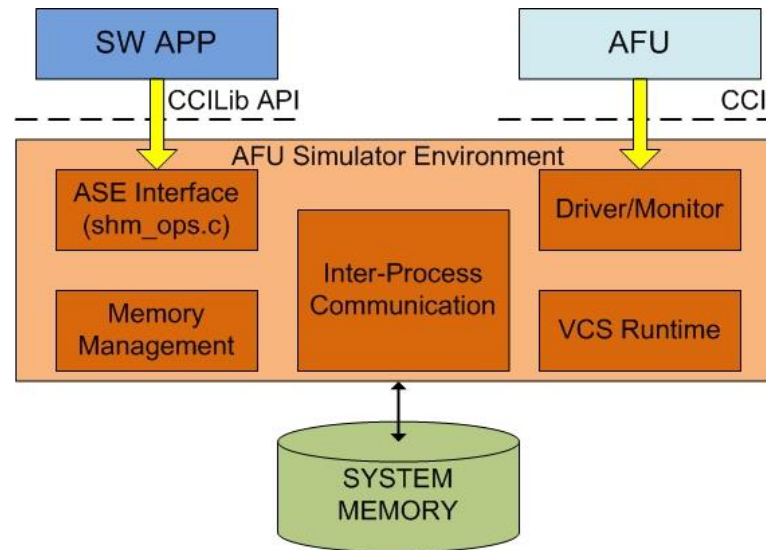
CA – QPI Caching Agent
 HA – QPI Home Agent
 QLP – FPGA implementation of QPI Link & Protocol Layer
 QPH – FPGA implementation of QPI Physical Layer
 MQ – Memory queue
 MC – Memory Controller

Pactron is QPI Licenses Provider

Intel QPI 1.1 Reference RTL Micro-Architecture

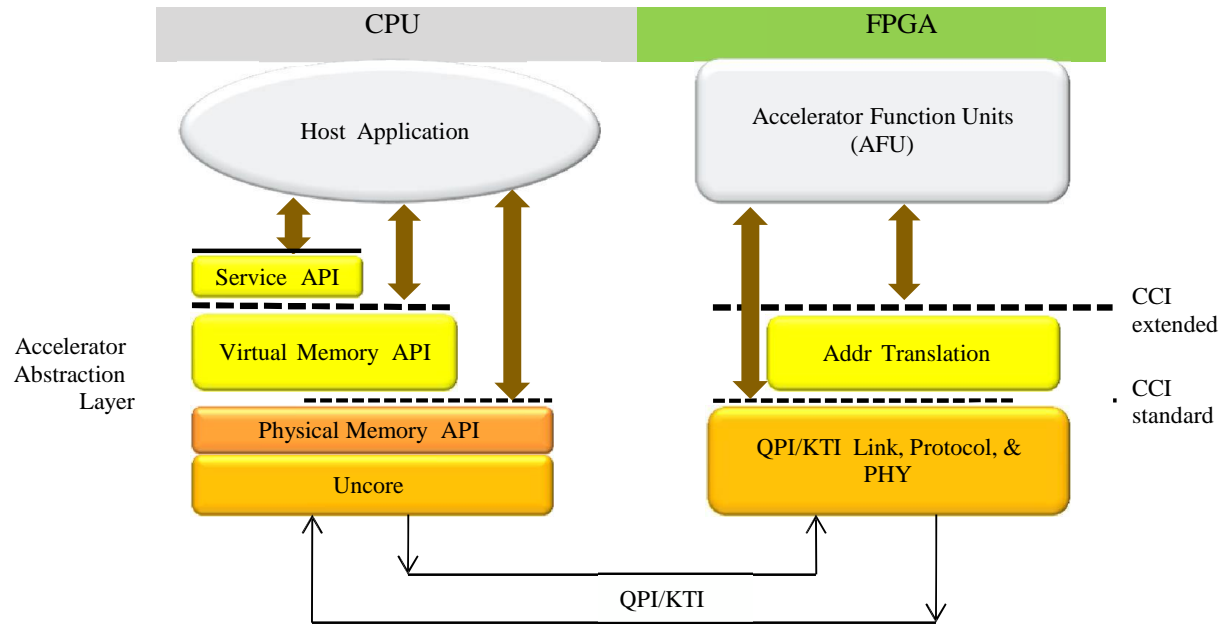


AFU Simulation Environment (ASE)



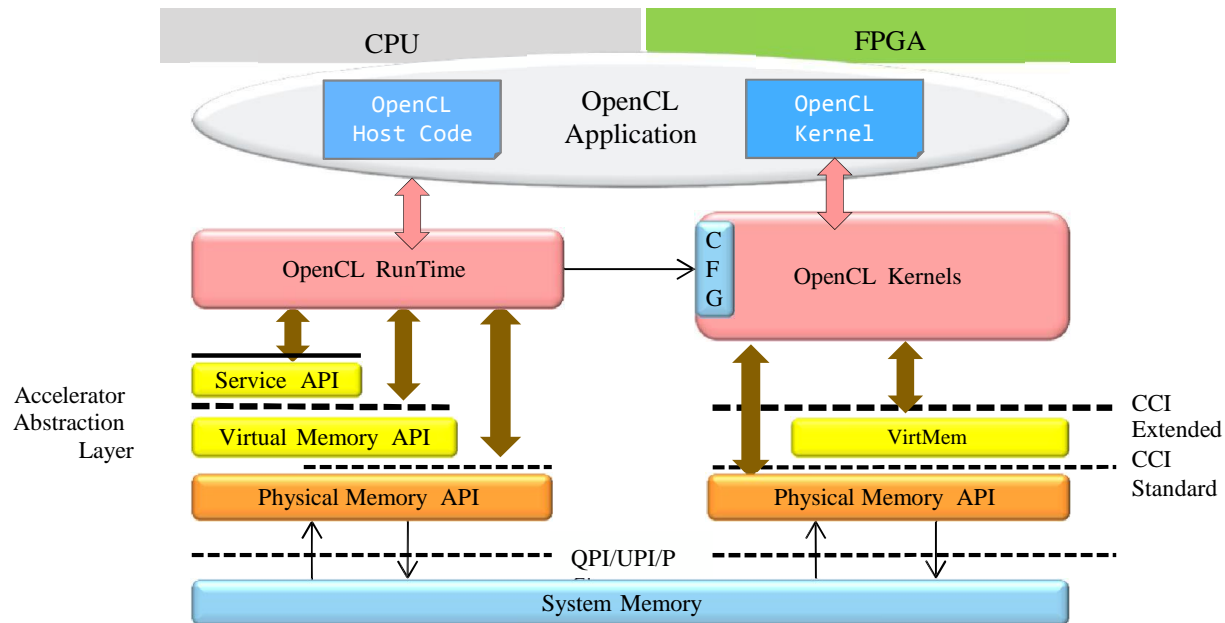
- Reduces hardware/software design cycle
- Allows end users to develop/test AFU RTL and software application in a single environment
 - Seamless portability from ASE development environment to Intel QuickAssist QPI-FPGA platform

Programming Interfaces: QuickAssist



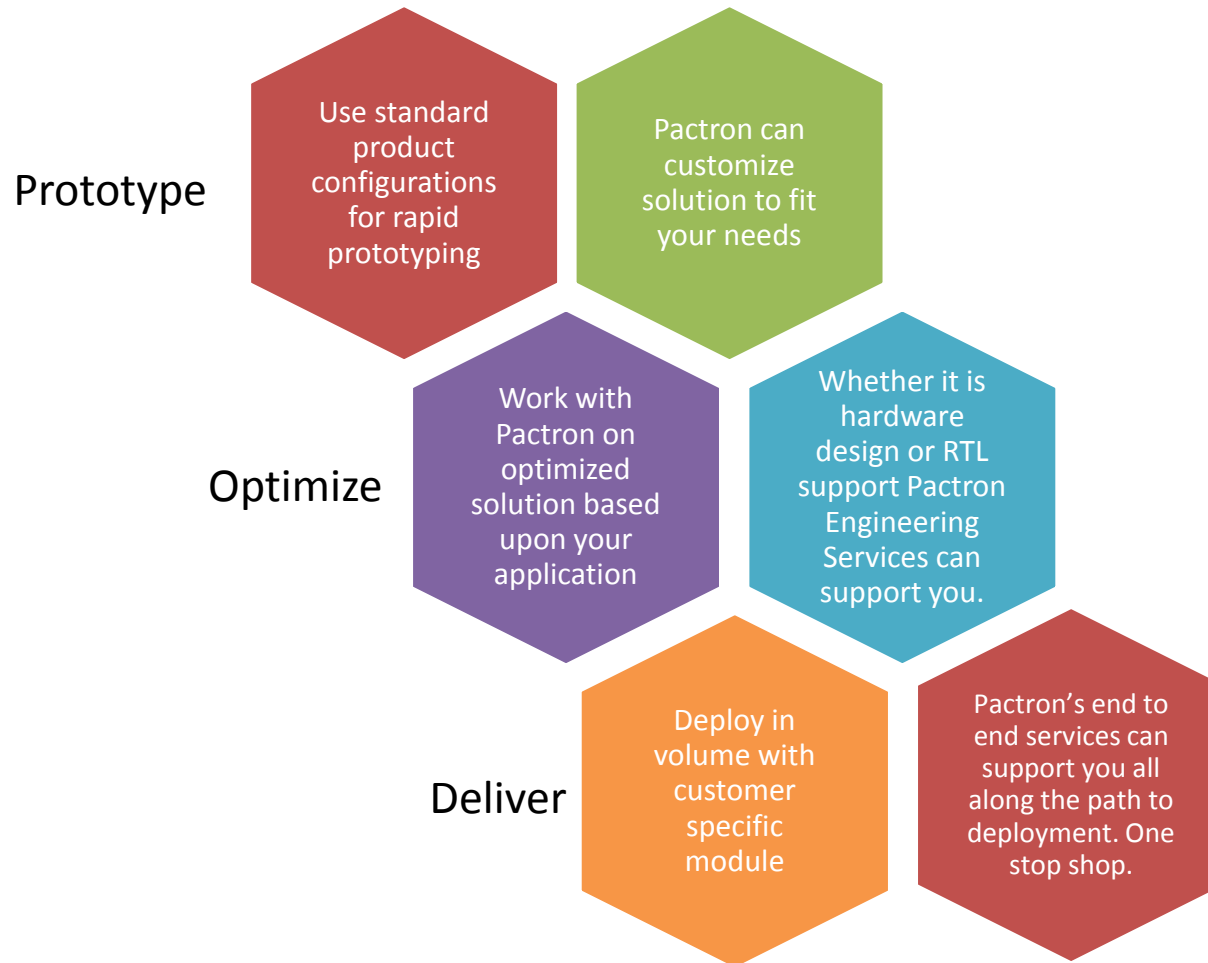
Programming interfaces will be forward compatible from SDP to future MCP solutions
Simulation Environment available for development of SW and RTL

Programming Interfaces ~ OpenCL



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

Pactron Integration Path



Pactron QPI Solutions Summary

- Intel® QPI Stack running on the Altera FPGA's
- Coherently attached to the shared memory space
- Cache inside the FPGA.....*this is a big deal!*
- Caching Agent and HA only with on-chip RAM
- Additional innovations are possible by merely reprogramming the Bitstream!
- **Work with Pactron to deliver a customized solution to fit your Application needs**