


High Performance Embedded Computing Module Enabled by GPU, CPU and FPGA SOCs



**Jason Fritz, Ph.D.,
Joshua Pearson and
Michael Bonato
Colorado Engineering, Inc.
September 11, 2013**



SBIR DATA RIGHTS

Contractor Name: Colorado Engineering Inc. (CEI)
Contractor Address: 1915 Jamboree Dr, Suite 165, Colorado Springs, CO 80920
Expiration of SBIR Data Rights: Expires 5 years after completion of project work for this or any follow-on SBIR contract, whichever is later.

This presentation contains data developed by Colorado Engineering under SBIR contract HQ0006-08-C-7908 and W31P4Q-12-C-0135. The Government's rights to use, modify, reproduce, release, perform, display, or disclose technical data or computer software marked with this legend are restricted during the period shown as provided in paragraph (b)(4) of the Rights in Noncommercial Technical Data and Computer Software - Small Business Innovation Research (SBIR) Program clause contained in the above identified contract. No restrictions apply after the expiration date shown above. Any reproduction of technical data, computer software, or portions thereof marked with this legend must also reproduce the markings.

Export or re-export of CEI products may be subject to restrictions and requirements of US export laws and regulations and may require advance authorization from the US Government.

The views expressed herein are those of the author and do not reflect the official policy or position of the Department of the Army, Department of Defense, or the U.S. Government. Reference herein to any specific commercial, private or public products, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement, recommendation, or favoring by the United States Government




Outline

- Background and Foundation
- 3D Processing
- The MPU/GPU/FPGA System on Chip (SOC) Module
 - MPU: AMD G-Series SOC (Goodbye bridges)
 - GPU: AMD E6760
 - FPGA: Altera Arria V SOC
 - PCIe mesh
- Conclusions


Background

- Ideal high performance computing environment is heterogeneous
 - Variety of processing technologies
 - General purpose / multi-core
 - Field Programmable Gate Arrays (FPGAs)
 - Graphics processors
 - Application Specific Integrated Circuits (ASICs)
 - **Enables engineer to optimize solution to fit Size, Weight, and Power (SWaP) budget**
- Technologies represent different points in trade space
 - Computing horsepower
 - Power consumption
 - Programmability



Background

- Heterogeneous computing approaches have traditionally relied on backplanes
 - Added weight, size, and cost
 - Constrain incremental scalability
- Backplanes limit designer's ability to realize solutions in physical volumes not conducive to legacy form factors
- VPX isn't VME – not as interchangeable and open



*Image courtesy of
Elma Bustronic*




Image courtesy of Kontron



A New Approach


- MDA funded an SBIR Phase I and II program to address challenges of modularity, scalability, and heterogeneity deployments with challenging form factors
 - RARE: Reconfigurable Advanced Rapid-prototyping Environment
 - Executed by CEI under technical guidance and influence of NRL, NSWC, and ONR
 - Recipient of 2011 Tibbetts Award
 - No Backplane!
- CEI and the Navy, sponsored by MDA, defined an open approach to SWaP-friendly embedded computing architectures
- **3DR** – RARE modules connected in 3D





3DR Modularity and Scalability

- 6.25" x 6.25" cards with interface connections in 3D
 - I/O bandwidth of 39 GB/sec per module via PCIe, LVDS, and SerDes
 - 3D direct connectivity of FPGA processing elements
 - I²C network of microcontrollers for health and status management
- Stack and/or tile modules in x, y, and z
 - Incrementally scale performance, I/O BW, and physical footprint
 - Physically reconfigure systems while maintaining common HW/FW/SW
 - Solutions in a fraction of the volume of traditional backplanes





COTS Modules: Digital and RF



GP/FPGA Processor

- AMCC 460SX PowerPC
- Xilinx Virtex-6 FPGA
- Dual 1Gb Ethernet
- USB, RS-232




ADC + FPGA

- 10 ADC channels
- 16b @ 160 MSPS
- Xilinx Virtex-6 FPGA



DAC + FPGA

- 2 DAC channels
- 16b @ 1GSPS
- Xilinx Virtex-6 FPGA



Dual 10Gb Ethernet



Phased Array Antenna / Radar Interface



Starter Kit



RF Up Converter




LO Synthesizer




High Fidelity, Low Phase Noise Clock Distribution

PCIe Expansion



2013 Roadmap


- Stratix/Arria V
- Multicore GP
- **GPU + x86 + FPGA SOCs**
- Ultra-wideband ADCs
- Enhanced Tamper Resistance



RF Down Converter




Fabric Communication without Dedicated Switch Cards



- PCIe switches built into modular architecture (Gen 3 for new boards)
- End points can be FPGAs or General Purpose Processors

- FPGAs also interconnect with low latency, high bandwidth across the 3D topology
 - LVDS
 - SerDes






Next Generation Connectivity

Supported I/O Bandwidth by connector type for Next Generation 3DR Modules

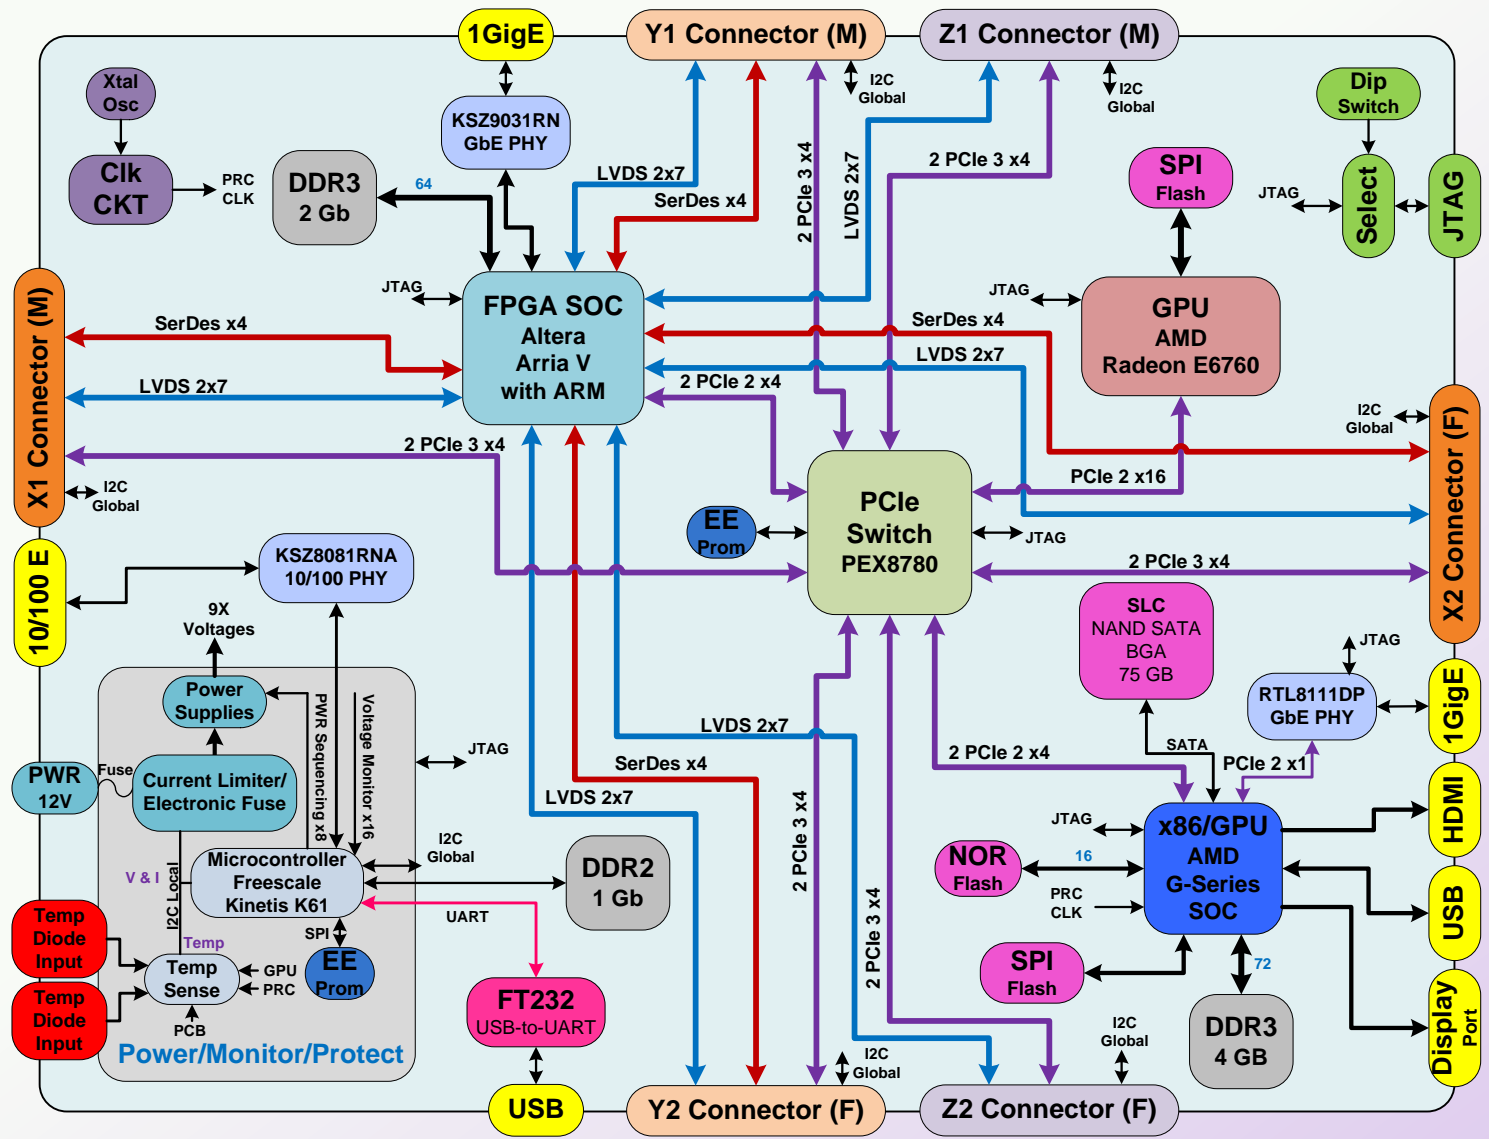
	1/2 Duplex	Full Duplex			Total Bandwidth (2 connectors) (Gb/s)
Connector	LVDS FPGA 14 pairs @ 1 GHz (Gb/s)	PCIe Gen 3 8 Lanes @ 8 GHz (Gb/s)	SerDes FPGA 4 Lanes @ 6 GHz (Gb/s)	SerDes PPC (Gb/s)	
X	14	64	24	1.25	210.5
Y	14	64	24	1.25	210.5
Z	14	64	24	3.125	214.25

Flex-3DR



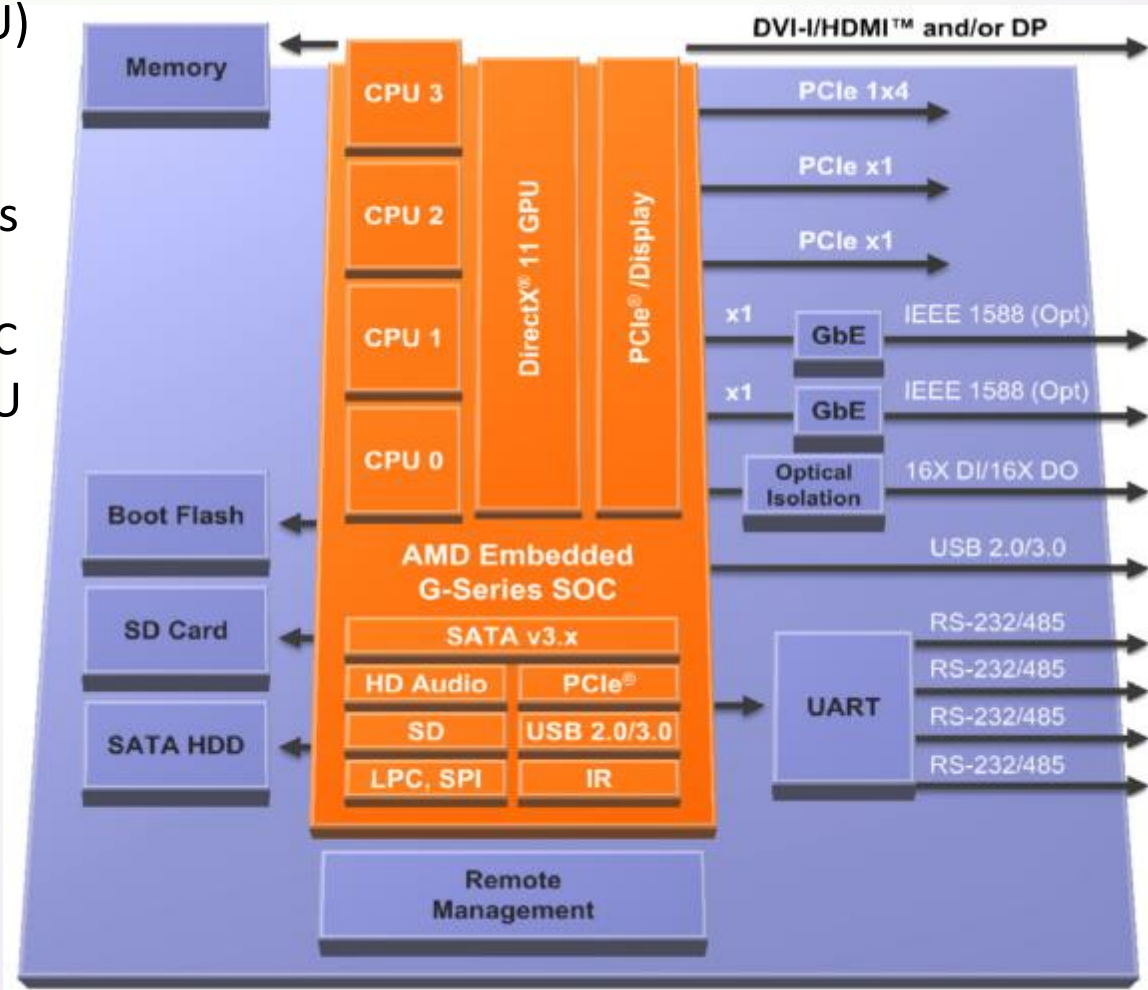


The GPU SoC Module



The Accelerated Processing Unit

- AMD G-Series SoC (MPU/GPU)
- North & South Bridge integrated on-chip
- Quad Core 2.0GHz MPU Cores
 - 64-bit data bus with 8-bit ECC
- Integrated Radeon 8400E GPU
 - 185 GFLOPS
- 25 W Thermal Design Power
- Gigabit Ethernet
- 75 GB NAND SATA
- 2 Mini Display Port
- USB 2.0



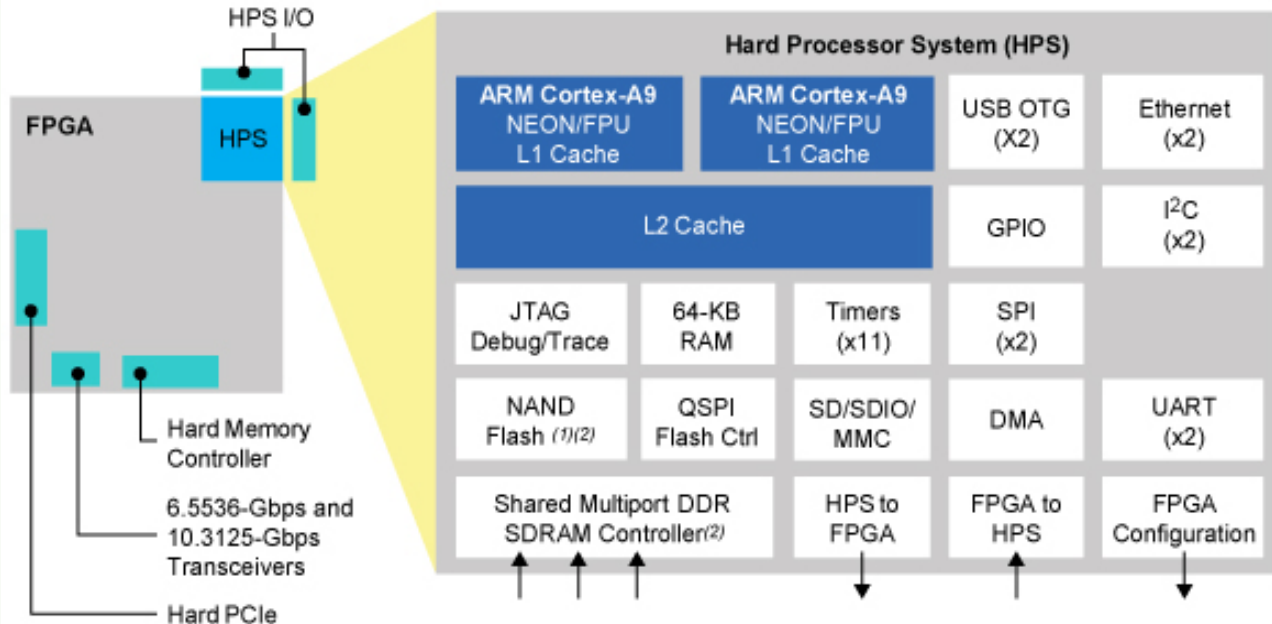


The Graphics Processing Unit

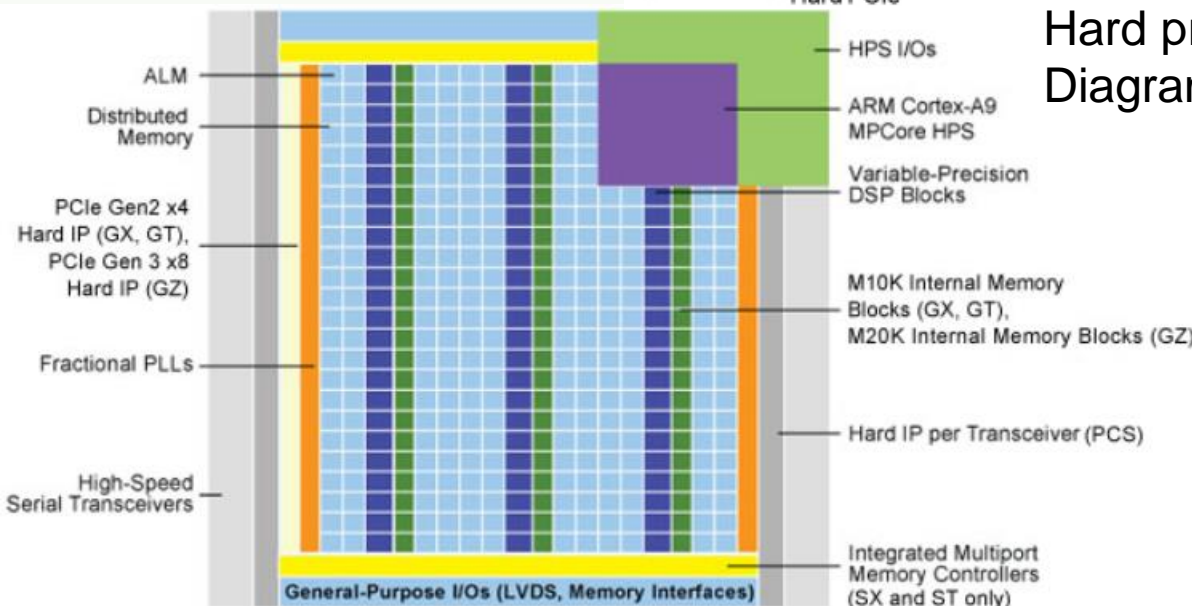
Feature	AMD Radeon™ E6760
Package Dimensions	GPU + memory, 37.5 x 37.5 mm BGA
Thermal Design Power (TDP)	35 W
Process Technology	40 nm
Graphics Engine Operating Frequency (max)	600 MHz
CPU Interface	PCI Express 2.0 (x1, x2, x4, x8, x16)
Shader Processing Units	6 SIMD engines x 80 processing elements = 480 shaders
Floating Point Performance (single precision, peak)	576 GFLOPS
Display Engine	AMD EyeSpeed visual acceleration, AMD Eyefinity, AMD HD3D
DirectX™ Capability	DirectX™ 11
OpenGL	OpenGL 4.1
Compute	AMD APP, OpenCL™ 1.1, DirectCompute 11
Internal Thermal Sensor	Yes
Memory	
Operating Frequency (max)	800 MHz / 3.2 Gbps
Configuration Type	128-bit wide, 1 GB, GDDR5

The FPGA with ARM Cores

- Altera Arria V with 2 integrated 800 MHz ARM Cortex-A9s
- 28 nm – low power



Hard processor System (HPS) Block Diagram



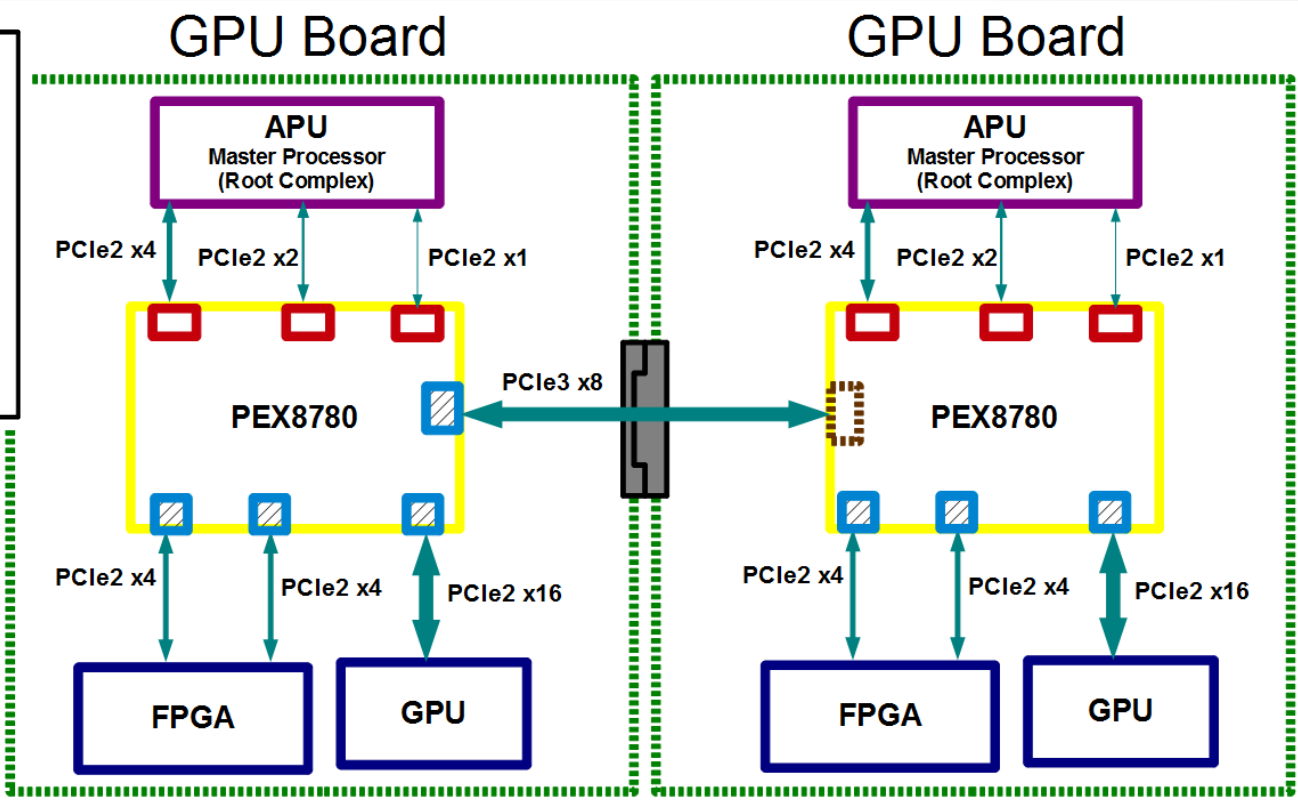

- 1600 GMACs
- 300 GFLOPs
- 125 Gbps between FPGA & HPS



Arria V SoC

FPGA Internal Memory	22.8 Mb
FPGA Logic Elements (LEs)	462,000
FPGA 18x18 Multipliers	2,136
FPGA 6.25 Gbps Transceivers	30
2 Hard PCIe Controllers	2 PCIe Gen3 x4 links
SERDES x4	SERDES x4 to 5 of 6 edge connectors
True LVDS	2 LVDS-7 to all 6 edge connectors
GMACS	1600 (FPGA + ARM Core)
GFLOPS	300 (FPGA + ARM Core)
Aggregate Bandwidth between ARM Core and FPGA	125 Gbps
ARM Core Off-board Peripherals	Gigabit Ethernet, UART-to-USB
ARM Cortex-9 Dual-Core Speed	800 MHz
FPGA Hard DDR3 Controller	256 MB DDR3
ARM Core Hard DDR3 Controller	256 MB DDR3

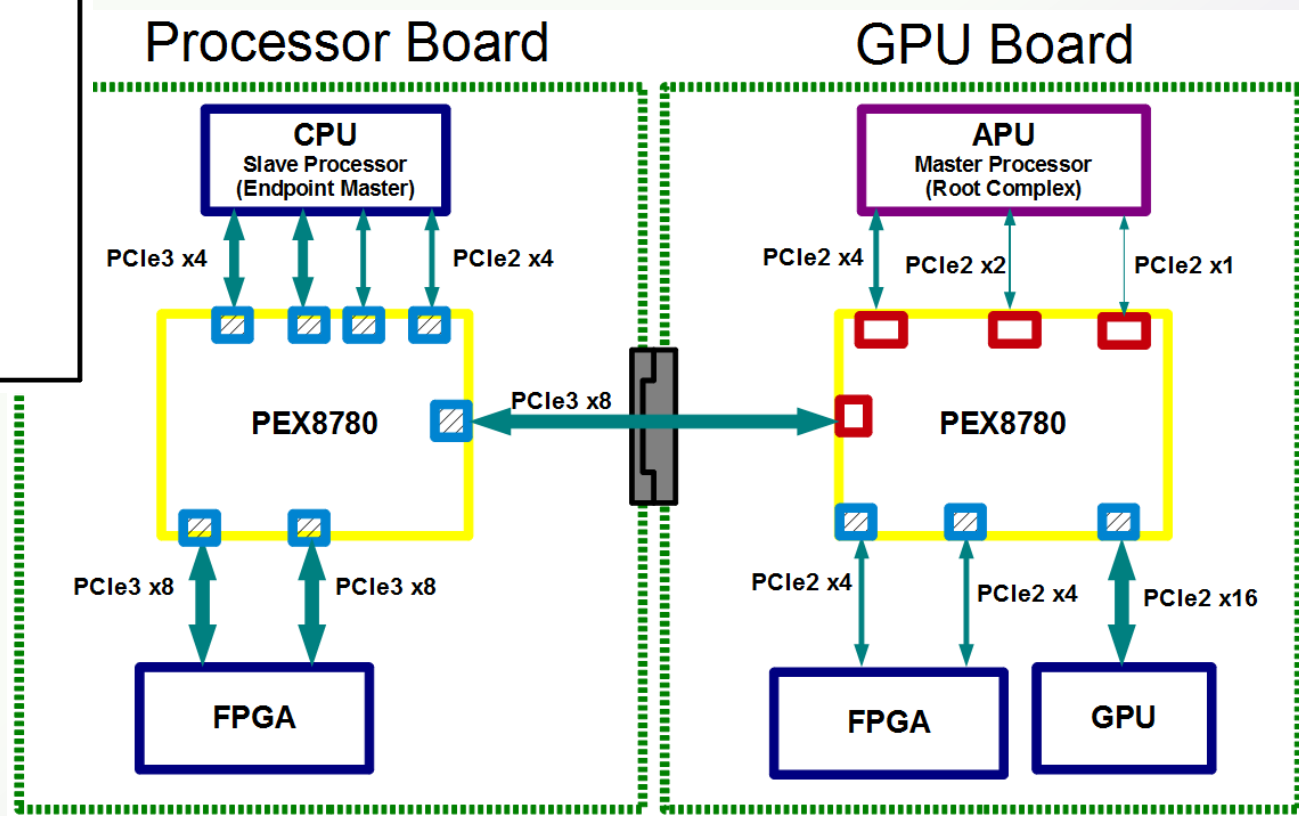
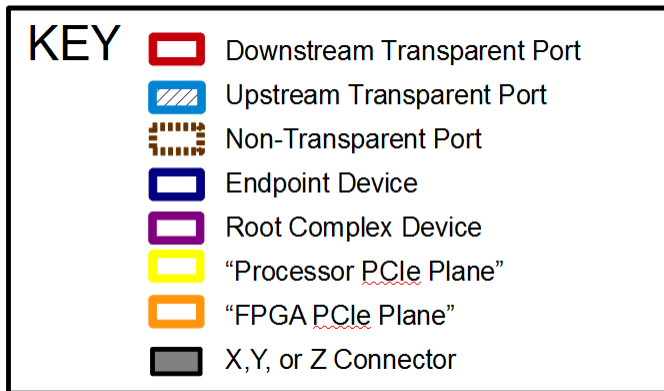
PCIe Mesh



- APU is root-complex only

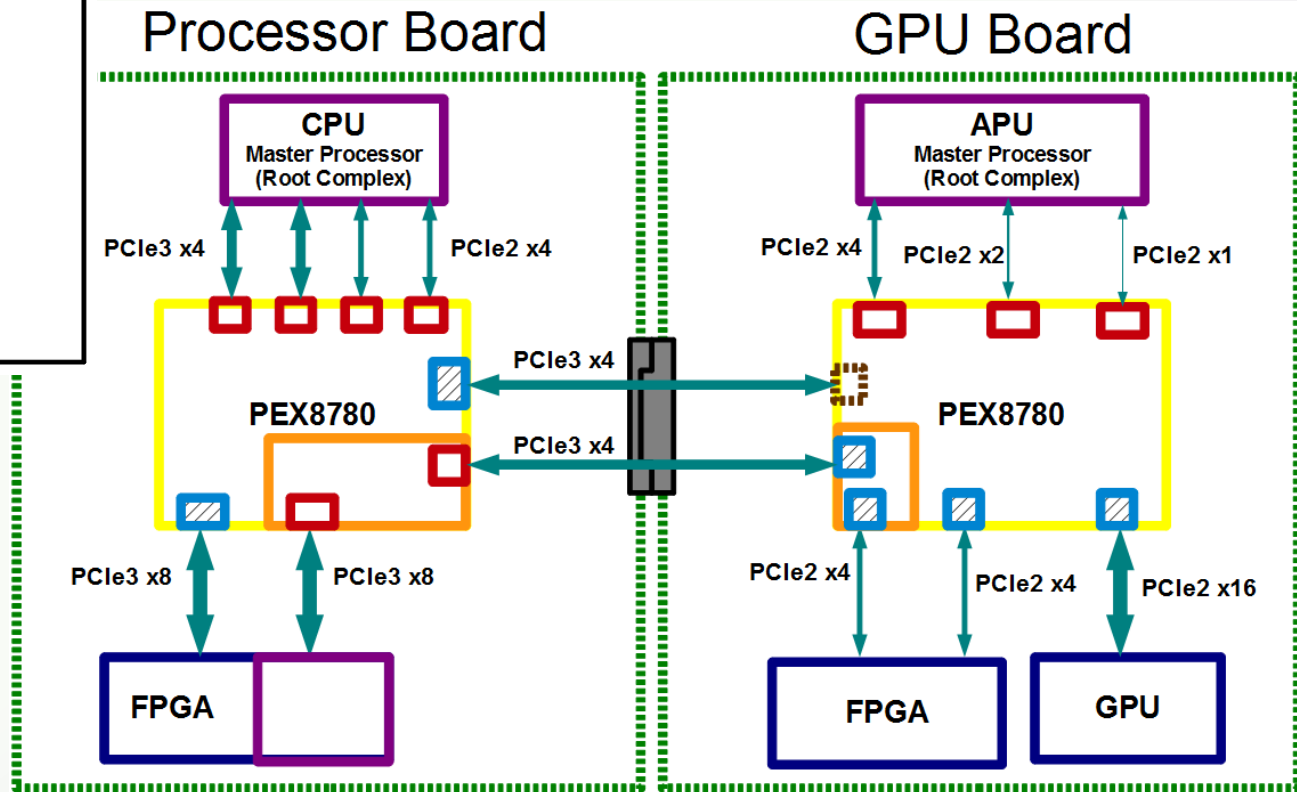
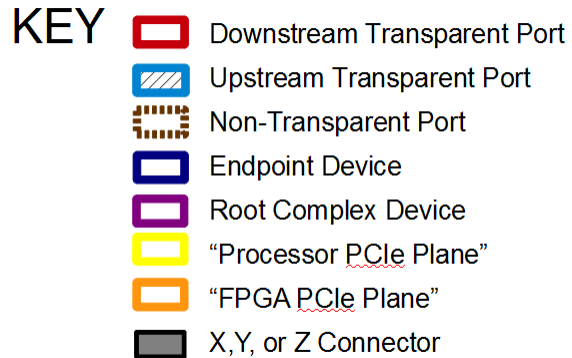
- Need to have one port configured to be Non-Transparent (NT)
 - During Enumeration, both APUs will treat the NT port as an endpoint from their domain
 - Address Translation Registers must be configured so that each Root Complex can access the PCIe memory address of the other domain
- Only one root complex will push/pull data to/from its domain

GPU to Processor Board



- CPU on Processor Board can be configured as an endpoint master
 - Eliminates the need for NT port... makes software a little more challenging

GPU to Processor Board



- Legacy CEI boards have 2 PCIe planes: Processor & FPGA
- PEX8780 can be configured as 2, 3, or 4 "virtual switches"
 - Each virtual switch is essentially an independent switch inside the PEX8780
 - 2 virtual switches creates 2 PCIe planes for backward X,Y,Z connectors supports 2 PCIe Gen3 x4 links or 1 PCIe Gen3 x8 link



Tying It All Together

- Heterogeneous compute platform on a single 6.25 in² PCB with connectors on all 6 sides
- AMD G-Series 4 MPU+GPU cores
- AMD E6760 GPU (576 GFLOPS)
- Altera Arria V w/ dual ARM Cortex-A9s
- PLX PCIe switch
- Provides tremendous flexibility to engineer at reasonable price for SWaP constrained applications
- Need more horsepower? Add more modules?
- Attach an ADC and DAC card for sensor processing
- Available Q1 2014
- Considering interface card for 2 MXM GPU modules



Thank You!

For more information please contact:

Jason Fritz, Ph.D.

jason.fritz@coloradoengineeringinc.com

719-388-8582 (main)

<http://www.coloradoengineeringinc.com>