

## OpenVPX Dual Cluster Architecture

### **Background: FPGAs achieve more speed per unit of power compared to CPUs, DSPs, and GPUs**

Advancements in silicon, software, and IP have proven FPGAs to be the ideal solution for accelerating applications on high-performance embedded computers and servers.

A study financed by the National Science Foundation (Alan George, Herman Lam, and Greg Stitt - IEEE magazine Computing in Science and Engineering - Jan/Feb 2011) provides a comparison of performance between reconfigurable-logic devices and fixed-logic devices. Figure 1 and 2 are illustrating this comparison in term of computational density in GOPS per Watt for 16 bit Integer and Single Precision Floating Point (SPFP) assuming an equal number of add and multiply operations.

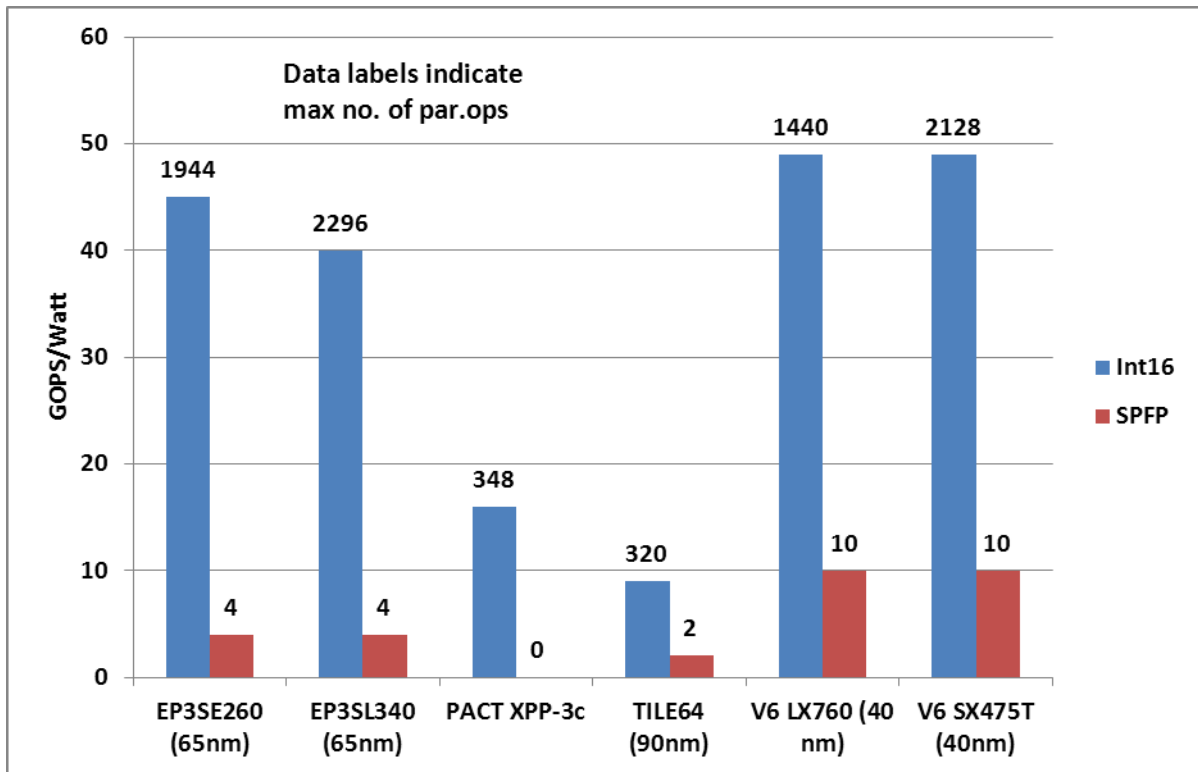


Figure 1: Reconfigurable-logic devices

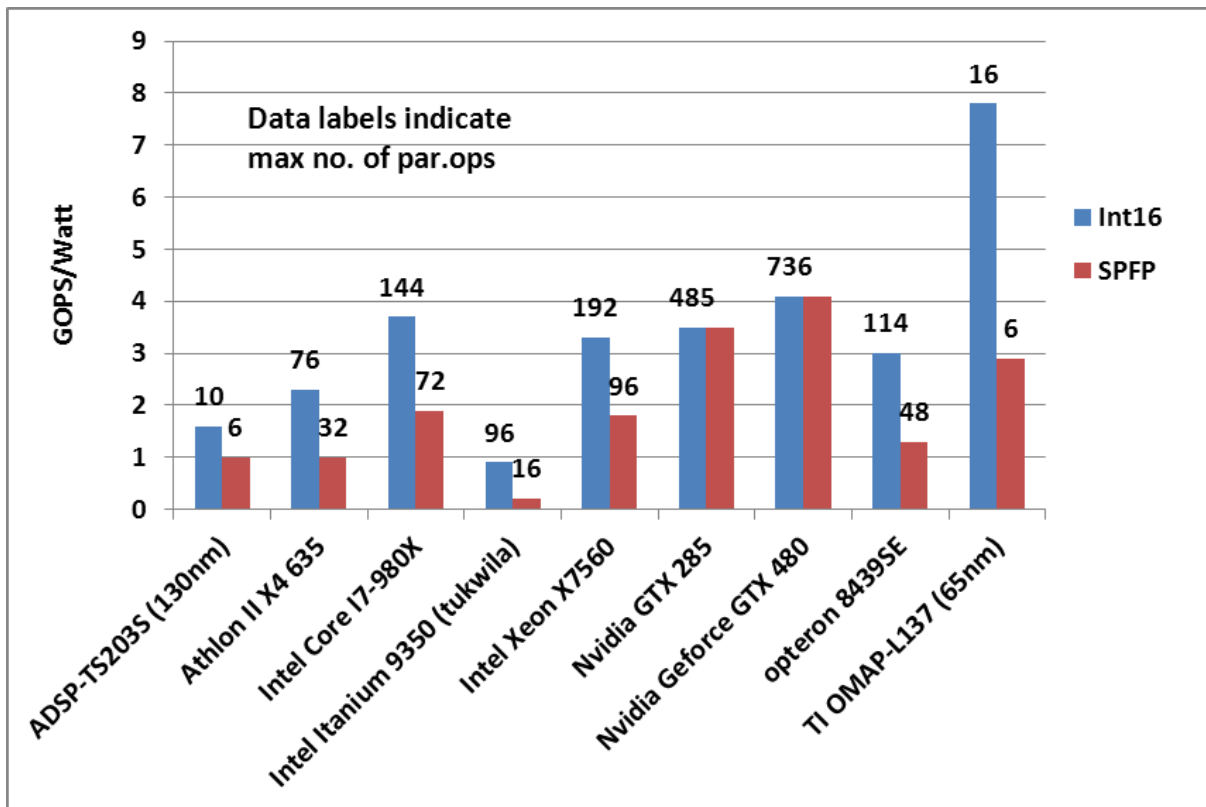


Figure 2: Fixed-logic devices

In general, FPGAs achieve more speed per unit of power compared to CPUs, DSPs, and GPUs for both types of operations. In these charts, the peak number of sustainable parallel operations of each type is cited atop each bar. These FPGA capabilities have opened up many opportunities to design advanced high performance computing platforms able to reach previously unattainable performance levels. These properties of the last generation of FPGAs are particularly useful for the design of the Digital Signal Processing embedded systems used in the military embedded Radar and Electronic Warfare applications where the power consumption must remain at low levels and where high computing performance is required.

## **New OpenVPX system architecture developed by Interface Concept.**

To take benefit from the high performance capabilities of the last generation of FPGAs, Interface Concept has developed a new processing Dual Cluster architecture. This architecture is characterized by a Signal Processing Cluster and a Data Processing Cluster as shown on Figure 3.

### **Signal Processing Cluster**

The Signal Processing Cluster consists in several FPGA resources and a control processor. This Cluster is responsible for initial data acquisition and preprocessing of the raw input data.

The role of the control processor can be to manage the FPGAs, to command and control the different sensors or for Double Precision Floating Point Computation.

Inside the Signal Processing Cluster, the general data communication between the FPGA Resources is made either through Fat Pipes PCIe x4 data paths across a hybrid OpenVPX switch (Data Plane) or through serial interconnects routed over the VPX Backplane (Expansion Plane). The data communication between the Processor and the other FPGAs is carried out through a Fat Pipe PCIe x4 data path across the hybrid switch. The control communication is carried out through the OpenVPX hybrid (\*) switch using Ultra-Thin Pipes 1000Base-X.

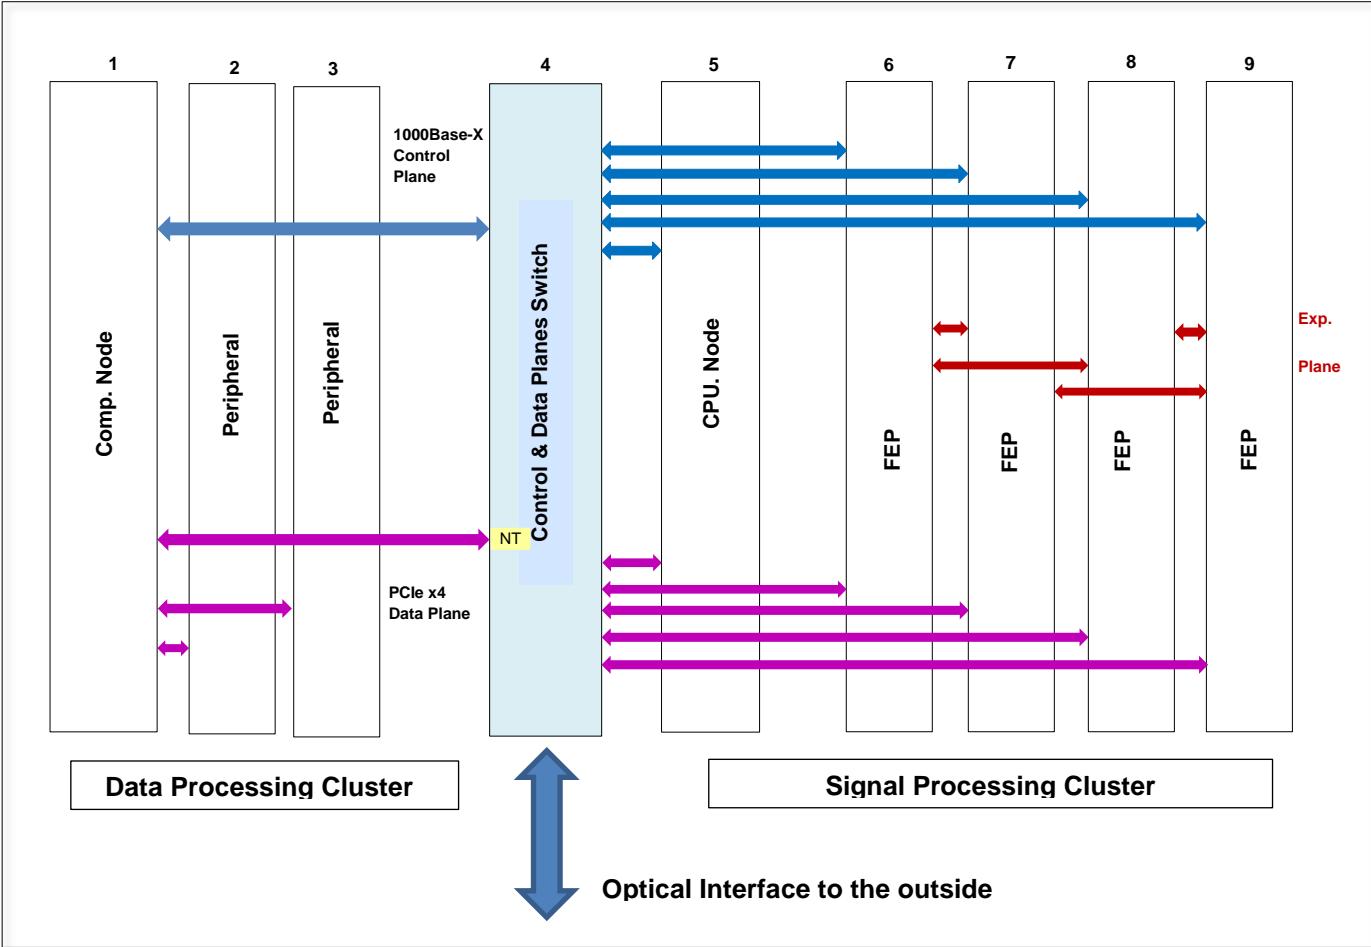
The data rate throughput of a Fat Pipe PCIe x 4 is according to the PCI Express 2.0 specifications. The typical data rate throughput of the Expansion Plane links using the Aurora 8b/10b protocol as an example, is up to 6.25 Gb/s per lane. The maximum throughput on each Ethernet link of the Control Plane is 1 Gbit/s.

### **Data Processing Cluster**

After being processed by the Signal Processing Cluster, the data can flow to the Data Processing Cluster for further processing. Post processing involves final computations and processes, or potentially a display or a storage process. The communication between the Signal and Data clusters are made across the hybrid switch both for the data paths and the control paths as shown in Figure 3. Inside this Data Processing Cluster the same communication principles apply as for the Signal Processing Cluster.

(\*) Hybrid means Dual Plane (Control and Data Plane)

The Non Transparent concept allows implementing multi-CPU architecture in the Data Processing Cluster area.

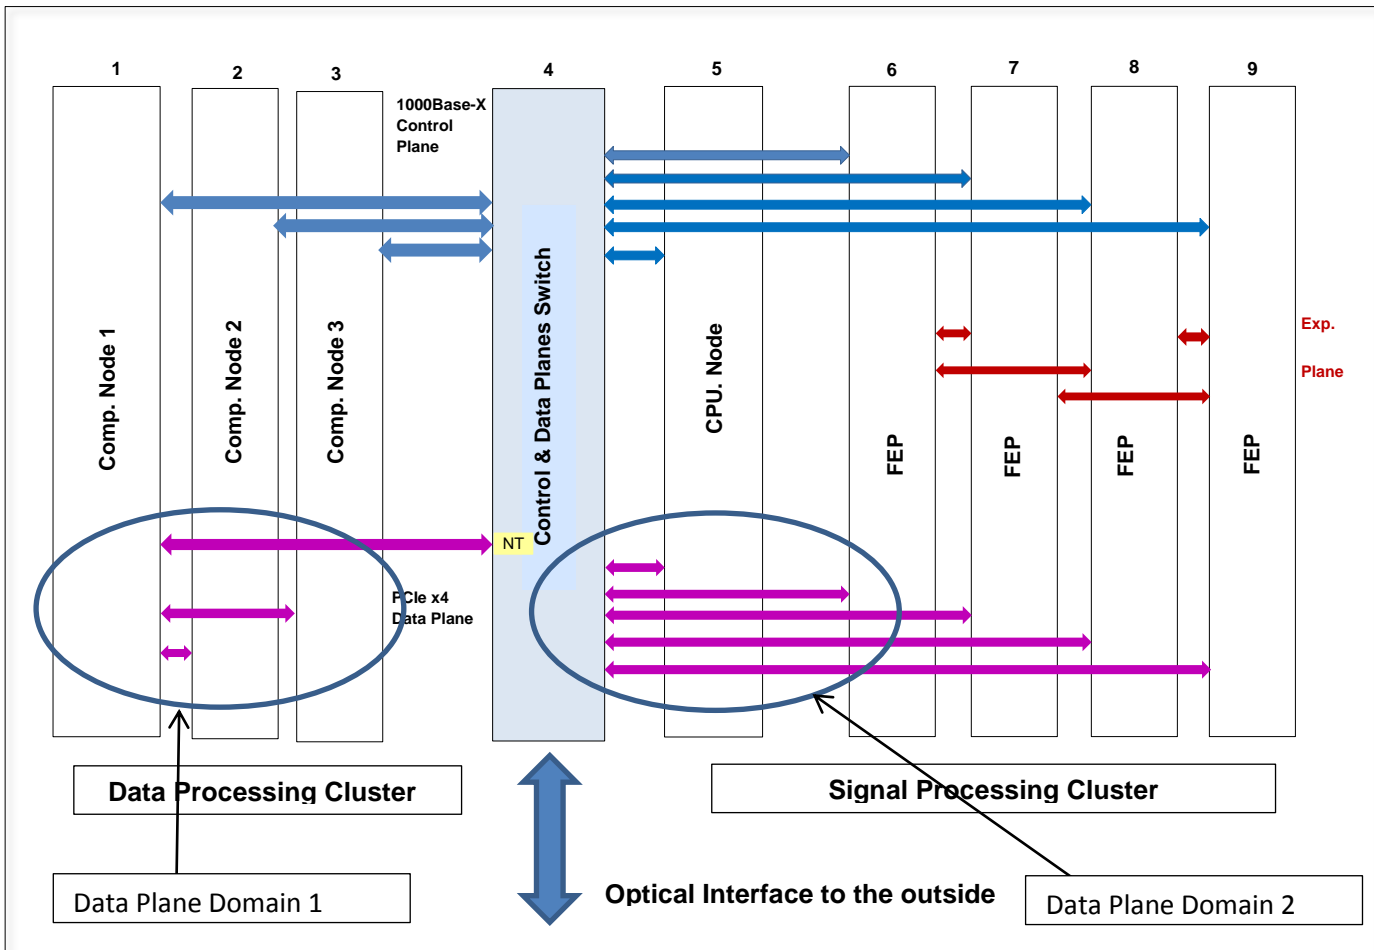


**Figure 3: Dual Cluster Architecture**

The hybrid switch (see Cometh 4410a below) can also provide on the Control Plane eight Ultra- thin Pipes allowing the connection of all the CPUs of the Data Processing Cluster to the Control Plane as shown in Figure 4.

**Cometh 4410a : a Hybrid OpenVPX Switch allowing low latency control and data interconnects**

The central piece of this architecture is the Hybrid OpenVPX Managed Low Power Cometh 4410a Switch. With its six PCIe ports each ensuring Fat Pipe communication, it allows up to 32 GBps of Bandwidth with low latency thanks to its cut-through architecture. Its Non-Transparent Bridging feature (NTB mode) allows the support of multiple NT Endpoints. It can therefore support several Root Complex. In its height 1000BASE-X ports configuration, it can interconnect all the resources of the Signal Processing and Data Processing Clusters on the Control Plane



**Figure 4: Dual Cluster Architecture with 3 Computer Nodes in the Data Processing Cluster**

providing full-wire speed L2 switching and L3 routing with L2-L4 advanced traffic classification, filtering and prioritization.

The Hybrid Switch is managed by a PowerPC Freescale processor running the powerful “Switchware” of Interface Concept. In addition one or two GbE ports allow communication to the outside world for the final output of the whole system for instance. The power consumption is below 15 Watts. This low power consumption allows the addition of an XMC directly connected to the Data Plane matrix of the switch as shown on the block diagram (Figure 5). This may allow saving one slot in the chassis. Adding an XMC on a high consuming power computer board may lead to serious difficulties in the temperature management.







**Figure 8: IC-FEP-VPX3b**

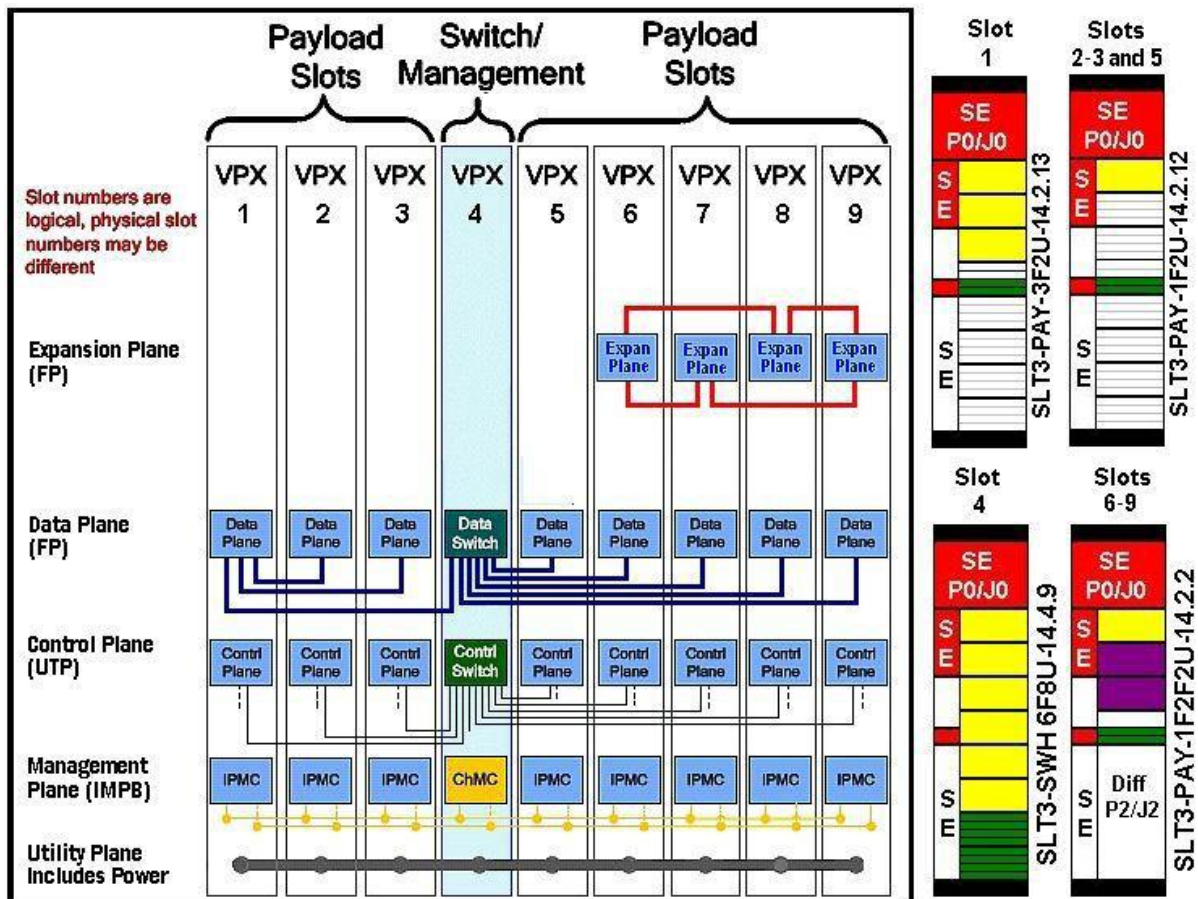
A reference design providing all the interface controllers, allows the developers to focus on the design of their processing algorithms. In addition a direct memory access (DMA) engine core with AXI4 interface can perform bulk transfers between two boards connected by PCI Express, reading data from a source address range and writing the data to a different address range. The DMA core has four channels, each using AXI4 protocol to connect AXI4 FPGA Peripherals.

It must be noted that the 6U version of this board named IC-FEP-VPX6a, can be considered as a compact version of this Dual Cluster Architecture. It features 2 interconnected Virtex-6 FPGAs each connected to two DDR3 40 bits wide banks and to an FMC and a QorIQ Freescale processor.



## Backplane

The Interface Concept backplane sustaining this co-processing Dual Cluster Architecture has been put into the OpenVPX VITA 65 norm under the Backplane Profile BKP3-CEN09-15.2.17-1 as shown in Figure 8.



**Figure 8: Co-Processing Dual Cluster Backplane BKP3-CEN09-15.2.17-1**

The following Interface Concept modules can be used:

- Slot 1 : IC-PPC-VPX3a 8640D), IC-PPC-VPX3b and IC-PPC-VPX3c with respectively QorIQ P3041 and P50x0 (available in Q1/2012), IC-DC2-VPX3a (SL9380) and the new Sandy Bridge Processing Unit (available in Q2/2012), IC-PQ3-VPX3a (8536E)
- Slots 2-3 : Idem as Slot 1 with the 6F8U version of the hybrid switch. In addition IC-GRA-VPX3a and IC-SSD-VPX3a
- Slot 4 : Cometh 4410a
- Slot 5 : Idem as Slot 1
- Slots 6-9 : IC-FEP-VPX3b and the future IC-FEP-VPX3i featuring a Virtex-7 FPGA

## Summary

Interface Concept has developed all the building blocks and payloads of a new high performance OpenVPX7 Dual Cluster system platform. The innovative architecture of this platform takes benefit from the high capabilities of the last generation of FPGA in terms of low power fast DSP processing that is required in the design of high performance Radar and EW systems for the military and aero markets.