


This Application Note describes a rugged, high performance radar processing unit which has been designed in conjunction with Cambridge Consultants Ltd. It combines NVIDIA® CUDA™ GPU processors with Virtex 6 FPGA and Intel CPU cards to create a scalable design. The system is built using VITA 46/48 (VPX) technology for use in the most rugged environments, and features a custom, multi-fabric backplane with PCI Express and Gbit Ethernet communication channels. The sealed chassis (conduction-cooled) uses mil-grade copper and fibre connections for high speed IO to the outside world. High capacity PSUs work in parallel to power multiple GPU/FPGA/CPU clusters.



Introduction

Modern Radar systems need to sustain real-time processing on multiple streams of high-bandwidth input data. The most efficient way of achieving this requires several types of processor working in parallel, linked by high capacity/low latency datapaths. At the front end, FPGAs are ideal for processing the incoming streams at full data rate (typically GBytes/s), using their massively parallel (but simple) signal processing blocks. Once the FPGA has reduced the data rate, the resulting stream is suitable for processing on a more general purpose platform. In the past this might have been a dedicated DSP processor (such as the TigerSHARC from Analog Devices) or a Freescale 8641D, but recently General Purpose Graphics Processor Units (GPUs) have become more popular. The latest GPU chips from ATI and NVIDIA contain hundreds of floating point processors, which can be controlled using standard C code, with some simple parallelisation syntax extensions.

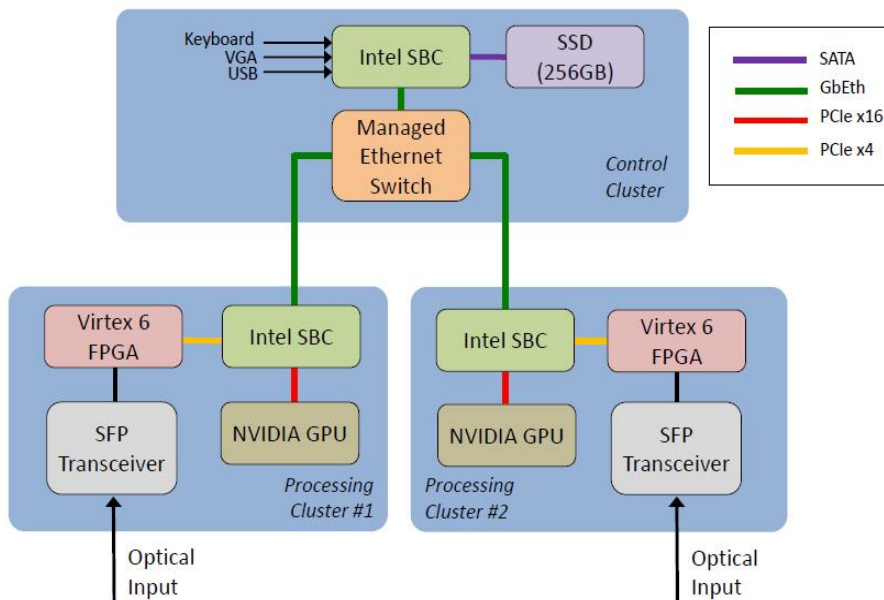


Figure 1 – System Architecture

Efficient communications between the processing elements is also critical – transfer times must be deterministic with low-latency. PCI Express has become a de facto standard which is supported in hardware in many modern processors (CPUs, FPGAs and GPUs). It can support dedicated point-to-point communications between processing elements, and also traditional PCI tree architectures for IO using PCIe switches. In this system the processing elements communicate via high speed serial channels using PCIe for data and GbEth for control (see Figure 1). The GbEth switch has spare capacity for controlling additional clusters and for providing a

simple external control/monitoring interface if required. Another feature of this system is component reuse to reduce the spares board count in deployment. The same SBC card is used as a cluster controller and a system controller which also simplifies the development by cutting the engineers' learning time.

System Description

The backbone of this system is a custom 3U VPX backplane designed to support PCIe channels between the cards. Unlike a traditional VME/cPCI chassis, VPX supports dedicated high speed serial channels between pairs of cards, allowing very high bandwidth transfers to happen in parallel. The backplane is housed in a 3U ATR-type chassis which can be cooled either via a cold wall, or liquid cooling plates attached to the sides of the chassis, depending on the deployment environment. All IO on the chassis uses sealed Glenair connectors, with mixed copper and fiber connections to the radar unit. The PSU is a scalable design, which can supply up to 750W into the chassis. Each PSU module is sealed and conduction cooled for use in harsh environments.

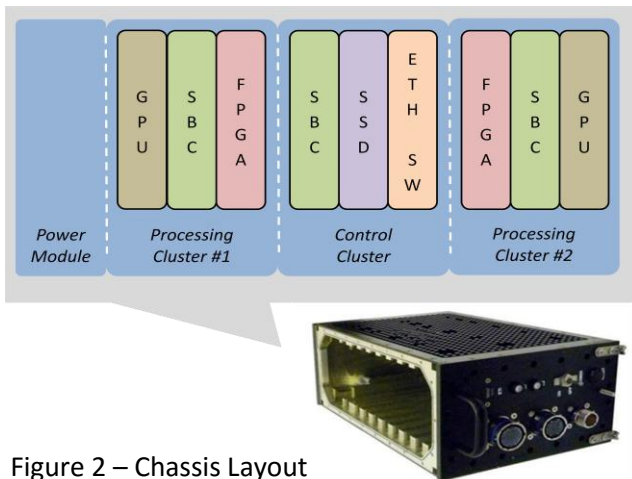


Figure 2 – Chassis Layout

The system uses a combination of FPGA (Virtex 6) and NVIDIA CUDA GPU processors, controlled via several Intel Core 2 Duo cards running Linux. The FPGAs are mounted on XMC modules which are fitted to VPX carrier cards, with all IO via the rear connectors (SFP optical transceivers are fitted to the backplane). The Virtex 6 FPGA supports PCIe endpoints in silicon, so this provides a simple way to interface the data stream between the FPGA and the CPU card. Optical data received by the FPGA is transferred via DMA into CPU memory in a ring buffer, and then sent over a separate independent PCIe channel from the CPU to the GPU board, via the backplane.

The system is built using 3U VPX cards from GE, one of the leading VPX card suppliers (range includes CPU & GPU cards, PMC/XMC carriers and Ethernet switches). Using cards from a single manufacturer substantially reduces the risks associated with integration.

- SBC341 – Intel Core 2 Duo processor coupled with up to 4GB RAM and a wide range of IO. Built in FLASH drive (4GBytes) to hold an OS or use the SATA channel for an external drive
- GRA111 – 3U VPX GPU card based on the NVIDIA GT240 CUDA processor. Dedicated PCIe x16 channel for high speed communication with SBC341
- GBX410 – Layer 2/3 Managed Ethernet switch with variety of front and rear IO options (up to 16 ports total). Also supports 10GbE XAUI ports to link multiple switches together
- PEX430 – XMC/PMC carrier for 3U VPX, incorporating a PCIe switch to simplify backplane configuration. In this application it hosts a custom Virtex 6 FPGA XMC module

Conclusion

This system design shows how VPX technology can be used to create a high bandwidth, high performance, rugged Radar platform, using multiple processor types (FPGAs and GPUs). FPGAs are used to process the high bandwidth incoming data streams, while GPUs perform more complex floating point operations on lower bandwidth data. The VPX backplane provides independent PCIe channels between cards for maximum system bandwidth. Choosing VPX modules from a single supplier reduces risk, but it is important that the system designer understands the power and cooling issues associated with such high power devices to ensure a reliable system.



AES Ltd are consultants and resellers of embedded products, targeting mil/aero and commercial applications in the UK.

For more information about AES products and services, please visit www.aes-eu.com

Advanced Embedded Solutions (AES) Ltd, Waverley House,

Lyndhurst Road, Christchurch, BH23 8LA

Tel +44 1425 517527

We would like to thank Cambridge Consultants Ltd for their input to this Application Note.