



Contribution ID: 79

Type: **Poster**

100Gbps PCI-express Readout for the LHCb Upgrade

Tuesday, September 23, 2014 5:03 PM (1 minute)

We present a new data acquisition system under development for the next upgrade of the LHCb experiment at CERN.

We focus in particular on the design of a new generation of readout boards, the PCIe40, and on the viability of PCI-express as an interconnect technology for high speed readout.

We show throughput measurements across the PCI-express bus in both directions, on Altera Stratix 5 devices, using a DMA mechanism and different synchronization schemes between the FPGA and the readout-unit.

Finally we discuss hardware and software design considerations necessary to achieve a throughput of 100Gbps per readout board.

Summary

A new requirement to allow a 'triggerless' readout of the entire detector at the LHC collision rate of 40MHz imposes a complete overhaul of the existing LHCb data acquisition. The new system will have to accommodate an aggregate bandwidth of several tens of terabits per second.

Designing such a system presents a compelling reason to study present and anticipated developments in interconnect technologies, not only from a performance perspective but also considering cost, maintainability and obsolescence.

This work explores the suitability of the 3rd generation of the PCI-express protocol for sustained, 100Gbps data acquisition workloads, in addition we present a data acquisition system based on specially designed PCI-express FPGA boards, called PCIe40.

Every PCIe40 aggregates 24 optical detector-readout channels and transmits processed event fragments into a readout computer over PCI-express. Readout computers are interconnected through a full-duplex local network where sparse event fragments are assembled into complete physics events.

We exploit modern FPGA devices (Altera Stratix 5 and Arria 10) where the PCI-express protocol is already integrated on-die as hardened logic, this directly translates to more programmable logic resources available for data processing.

We describe a streaming, high-performance DMA controller that was specifically implemented for the PCIe40 and we show how it integrates in the overall dataflow of the generic readout firmware being developed in parallel.

As the requirements for LHCb are very close to the practical limits of current PCI-express technology, this study will also provide insights on DMA (Direct Memory Access) performance measurement and optimization that will hopefully prove useful for other data acquisition scenarios outside of our particular use case.

In particular, given the tight coupling between hardware and software in such a system, we discuss how both the digital logic and the driver software are designed in conjunction to maximize overall readout performance.

Although PCI-express exhibits a number of interesting technical challenges, we show how such an implementation can result in a data acquisition system which is compact, economically advantageous and able to satisfy its design requirements.

The final system is scheduled to be deployed at the LHCb experiment in 2020.

Primary authors: BALBI, Gabriele (INFN); NEUFELD, Niko (CERN); DURANTE, Paolo (CERN); SCHWEMMER, Rainer (CERN); MARCONI, Umberto (INFN)

Presenter: DURANTE, Paolo (CERN)

Session Classification: First Poster Session

Track Classification: Logic