



Contribution ID: 202

Type: Poster

## FPGA implementation of PCI-express bifurcation for high-throughput data acquisition

*Tuesday 29 September 2015 17:46 (1 minute)*

The experiments at the LHC are undergoing a massive design upgrade to increase their data-taking capacities in the coming years, in anticipation of higher luminosity and new running conditions. For some experiments, the requirement of 100Gbps of readout bandwidth per readout unit has driven the adoption of PCI-express Gen3 as the main readout protocol. Limitations of current FPGA silicon to 8 lanes per interface require exploiting multiple interfaces to achieve the desired performance. In this work we study how PCI-express lane bifurcation could be exploited to overcome this limitation while minimizing BOM and layout complexity.

### Summary

The LHCb experiment at CERN is due to be upgraded to a purely software trigger in the coming years, this requires a complete redesign of the entire readout chain in order to achieve a continuous ('triggerless') readout of all subdetectors at the LHC collision rate of 40 MHz.

As part of this upgrade, a new readout board, common to all LHCb subsystems, is currently under development at CPPM (Center for Particle Physics of Marseilles). PCI-express Gen3 has been designated as the main communication protocol between the on-board FPGA performing the initial data processing and the CPU running the distributed event-building algorithm.

Current FPGAs already implement the PCI-express protocol in hardened logic, however these implementations are limited to 8-lane interfaces. To maximize the efficiency of the event-building network, the throughput from the readout board should match the throughput available on the event-building network, this requires the use of two PCIe interfaces on each readout FPGA.

Since PCI-express is a point-to-point protocol, one way to connect the CPU with these interfaces is by adding a dedicated hardware switch to the readout board. In order to avoid the additional cost, power and routing complexity of this approach, we study an alternative solution exploiting PCI-express lane bifurcation.

Lane bifurcation refers to a design feature of modern CPUs where a single PCI-express root port can be reconfigured to appear as two or more logical root ports, each using a subset of the original communication wires.

After validating the compatibility of bifurcation with Altera FPGAs, using a commercial development kit, we examine the mechanism required to configure bifurcation on a variety of server platforms, both through cooperation with BIOS vendors and through direct low-level I/O.

For the latter approach, after initial prototyping inside the Intel BITS environment to perform low-level access to the CPU, we turn to implement a mechanism able to execute the same operations directly from the FPGA board.

For this purpose, the interface to the PCI-express HardIP on the FPGA is modified in order to expose an embedded ROM to the server BIOS, in accordance with the BIOS PnP specification. Successively, the ROM is programmed with low-level, executable code that enumerates the PCI hierarchy and enables bifurcation on the appropriate root port.

Lastly, this implementation is ported to an experimental Arria 10 development kit in order to measure the combined data acquisition throughput that can be obtained through the bifurcated interfaces.

**Primary author:** DURANTE, Paolo (CERN)

**Co-authors:** NEUFELD, Niko (CERN); SCHWEMMER, Rainer (CERN)

**Presenter:** DURANTE, Paolo (CERN)

**Session Classification:** Poster

**Track Classification:** Logic