



Contribution ID: 192

Type: **Poster presentation**

Particle identification on an FPGA accelerated compute platform for the LHCb Upgrade.

Tuesday 7 June 2016 15:00 (1h 30m)

The current LHCb readout system will be upgraded in 2018 to a 'triggerless' readout of the entire detector at the LHC collision rate of 40 MHz. The corresponding bandwidth from the detector down to the foreseen dedicated computing farm (event filter farm), which acts as the trigger, has to be increased by a factor of almost 100 from currently 500 GBit/s up to 40 TBit/s. The event filter farm will pre-analyse the data and will select the events on an event by event basis. This will reduce the bandwidth down to a manageable size to write the interesting physics data to tape.

The design of such a system is a challenging task, why different technologies are considered and have to be investigated for the different parts of the system. For the usage in the event building or in the event filter farm (trigger) an experimental FPGA accelerated computing platform is considered and therefore tested. FPGA compute accelerators are more and more used in standard servers like for Microsoft Bing search or Baidu search. The platform we use hosts a general CPU and an high performance FPGA linked via an high speed link. On the FPGA an accelerator is implemented. The used system is a two socket platform from Intel with a Xeon CPU and an FPGA. The CPU and the FPGA are connected via the point-to-point interconnect QPI, which is used to interconnect CPUs in industry standard server. The FPGA has cache-coherent memory access to the main memory of the server and can collaborate with the CPU. These cache-coherent architectures are better suited for real-time connections between FPGA and CPU as the usual PCIe FPGA accelerators. It is very likely that these platforms, which are built in general for high performance computing, are also very interesting for the High Energy Physics community.

As First step it is tested to port the existing LHCb RICH particle identification to the experimental FPGA accelerated platform. We will compare the performance of the LHCb RICH particle identification running on a normal CPU with the performance of the same algorithm, which is running on the Xeon-FPGA compute accelerator platform. Furthermore, the performance results of smaller test cases performed at the beginning like sorting are presented.

This work is done in collaboration with Intel Corporation.

Author: FAERBER, Christian (CERN)

Co-authors: Mr MACHEN, Jonathan (Intel); NEUFELD, Niko (CERN); SCHWEMMER, Rainer (CERN)

Presenter: FAERBER, Christian (CERN)

Session Classification: Poster session 1