Contribution ID: **122** Type: **Poster**

A High-Speed DAQ Framework for Future High-Level Trigger and Event Building Clusters

Wednesday, 28 September 2016 18:18 (1 minute)

Modern data acquisition and trigger systems require a throughput of several GB/s and latencies down to the microsecond level. In order to satisfy such requirements, we developed a heterogeneous system with FPGA-based readout cards and GPU-based computing nodes coupled by fast links. Remote DMA engines are used for direct communication between Xilinx FPGAs and GPUs from AMD / "DirectGMA" and NVIDIA / "GPUDirect". Scalability is ensured by InfiniBand interconnects using the same technologies. In this contribute we present the system architecture and we compare the performance of the different solutions in terms of data throughput and latency.

Summary

Significant new challenges are continuously arising in High Energy Physics (HEP) experiments at the Large Hadron Collider (LHC) at CERN. The quest for rare new physics phenomena leads to the evaluation of Graphics Processing Units (GPU) as enhancement for the existing high-level trigger (HLT). The trigger upgrade requires faster and more efficient event selection. GPUs with its flexibility might also induce the possibility of new complex triggers that were not feasible previously. A HLT with efficient many-core parallelization of the event reconstruction offers the opportunity to significantly reduce the number computing nodes. To obtain the best results, HPC clusters with multiple GPUs require high-performance interconnect such as InfiniBand to handle the GPU-to-GPU communications and to optimize the overall performance of the cluster. Because the GPUs put significant demands on the interconnects, high-performance links like InfiniBand are required to provide low latency, high message rates, and bandwidth in order to enable all resources in the cluster to run at peak performance.

The Institute for Data Processing and Electronics (IPE) at the Karlsruhe Institute of Technology has developed an integrated data acquisition framework that provides on-line data processing in the range of GB/s. The integration of GPUs in trigger and data acquisition systems is for example used in HEP and synchrotron experiments. To optimize data throughput and latency, we have developed a custom readout architecture that enables fast communication between FPGA, GPUs and InfiniBand via Remote Direct Memory Access (RDMA). In this contribute we present the high throughput platform that is capable to transfer data from FPGA to GPUs and FPGA to InfiniBand by RDMA technologies. The proposed architecture is a candidate for future generations of event building clusters to collect all event fragments belonging to the same L1 trigger and to transmit the final events to the nodes of the HLT filter farms. Direct GPU connection could be used for high-performance HLTs where GPUs exchange data with GPU by InfiniBand or FPGA for a low latency processing.

Primary author: CASELLE, Michele (KIT - Karlsruhe Institute of Technology (DE))

Co-authors: Dr KOPMANN, Andreas (KIT - Karlsruhe Institute of Technology (DE)); MOHR, Hannes (KIT - Karlsruhe Institute of Technology (DE)); ROTA, Lorenzo (KIT - Karlsruhe Institute of Technology (DE)); E.

ARDILA, Luis (KIT - Karlsruhe Institute of Technology (DE)); WEBER, Marc (KIT - Karlsruhe Institute of Technology (DE)); BALZER, Matthias Norbert (KIT - Karlsruhe Institute of Technology (DE)); VOGELGESANG, Matthias (KIT - Karlsruhe Institute of Technology (DE)); CHILINGARYAN, Suren (KIT - Karlsruhe Institute of Technology (DE)); Mr DRITSCHLER, Timo (KIT - Karlsruhe Institute of Technology (DE))

Presenter: CASELLE, Michele (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: POSTER

Track Classification: Trigger