



Contribution ID: 99

Type: Oral presentation

High-speed, low-latency readout system with real-time trigger based on GPUs

Friday, June 10, 2016 9:10 AM (20 minutes)

Significant new challenges are continuously confronting the 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 a Graphics Processing Unit (GPU) enhancement for the existing high-level trigger (HLT), made possible by the current flexibility of the trigger system, which not only provides faster and more efficient event selection, but also includes the possibility of new complex triggers that were not previously feasible.

At HLT, when the efficient many-core parallelization of event reconstruction algorithms is possible, the benefit of significantly reducing the number of the farm computing nodes is evident.

At lower levels, where typically severe real-time constraints are present, we envisioned the possibility to meet the real-time constraints and to reduce data transfer latency and its fluctuations, by injecting readout data directly from the FPGA into the GPU memories without any intermediate buffering, therefore offloading the CPU, avoiding OS jitter effects.

In order to satisfy such constraints at lower levels, we have developed a custom FPGA-based readout card and implemented a new concept of Direct Memory Access (DMA) capable to move the data from FPGA to system memory and/or GPU memory. The readout card is equipped with a Xilinx Virtex-7 FPGA and it is connected to a GPU farm by a generation 3 PCIe x16 data link, capable of a net throughput of up to 13 GB/s. We have integrated the DMA engine with AMD's "Direct GMA" technology to enable data transfers to GPU memory with a measured data throughput of up to 6.4 GB/s in x8 lanes operation mode.

For GPU algorithm, a tracking algorithm for transverse momentum p_T trigger is evaluated on a NVIDIA Tesla K40 GPU using Hough-transform methods. A prominent result shows that 500 Stubs are elaborated in only 13 μ s with only one GPU core. These results show that low GPU elaboration times combined with low latency and high throughput electronics open a new prospective for a GPU-based low-level trigger system for the CMS experiment.

Benchmarks for latency and bandwidth for the proposed readout system are presented, followed by a performance analysis on case studies of the GPU-based low level trigger for the CMS experiment. In addition, the use of DMA in the form of NVIDIA's "GPU Direct" and InfiniBand for low-level trigger will be discussed. Finally, we give an outline of future project activities.

Primary author: Dr CASELLE, Michele (Karlsruhe Institute of Technology)

Co-authors: Dr KOPMANN, Andreas (Karlsruhe Institute of Technology); Mr MOHR, Hannes (Karlsruhe Institute of Technology); ROTA, Lorenzo (Karlsruhe Institute of Technology); Mr ARDILA PEREZ, Luis Eduardo (Karlsruhe Institute of Technology); WEBER, Marc (KIT - Karlsruhe Institute of Technology (DE)); BALZER, Matthias Norbert (KIT - Karlsruhe Institute of Technology (DE)); VOGELGESANG, Matthias (Karlsruhe Institute of Technology); CHILINGARYAN, Suren (Karlsruhe Institute of Technology); Mr DRITSCHLER, Timo (Karlsruhe Institute of Technology)

Presenter: Dr CASELLE, Michele (Karlsruhe Institute of Technology)

Session Classification: Trigger 2

Track Classification: Trigger Systems