



Contribution ID: 194

Type: Poster

A multi-Gigabyte per Second PCI-Express Data Link for Real-Time DAQ Systems

Tuesday, September 29, 2015 5:45 PM (1 minute)

A growing number of physics experiments requires DAQ systems with multi-Gbytes/s data-links. We developed a Direct Memory Access (DMA) engine compatible with the Xilinx PCI Express Gen2/3 core. The design comes with a Linux driver. Preliminary measurements with a Gen3 single-core have reached a throughput of up to 6.7 GBytes/s. We also intend to use this technology for direct communication between FPGA-based DAQ electronics and GPU memories (e.g. NVIDIA GPU-direct). This architecture finds its application in real-time DAQ systems and in low and high-level triggers for HEP experiments.

Summary

Recent years have seen an increasing interest in GPU-based systems for HEP applications, which require high data rates and high computational power (e.g.: ATLAS, CMS, ALICE, Mu3e and PANDA high/low-level triggers). Moreover, the volumes of data produced in recent photon science facilities have become comparable to those traditionally associated with HEP. In such experiments data is acquired by one or more read-out boards and then transmitted to high-end external processing units in short bursts or in a continuous streaming mode. With expected data rates of several Gbytes/s, the data transmission link between the read-out boards and the host system can constitute the performance bottleneck.

To address this problem we have developed an architecture that connects FPGA-based devices and external processing units by PCIe data links. The Direct Memory Access (DMA) engine is also fully compatible with the GPU-direct technology: instead of accessing the central CPU memory, the engine directly transfers data in the GPU internal memory. By bypassing the CPU, the transmission latency and the memory bandwidth requirement of the system are reduced.

The high-performance and compact architecture is fully compatible with the Xilinx PCIe Gen2/3 core for the FPGA families 6 and 7, and its integration with a custom FPGA design is done with minimum effort. The hardware engine is interfaced via a custom-designed Linux driver.

Our implementation has been optimized in order to achieve the maximum data throughput and to minimize the FPGA resource utilization, while still maintaining the flexibility of a Scatter-Gather memory policy. The architecture includes a simple and configurable Base Address Registers (BAR) access, which is used to program the DMA engine and the external application-specific logic. Data is provided to the DMA engine via a user-friendly FIFO interface, with a data word width of 128 bits for Gen2 and 256 bits for Gen3 (according to the input/output data width of the PCIe core), operating at 250 MHz. A dual-core engine has been realized by implementing two PCIe Gen2 x8 cores in parallel and connecting them to an external x16 PCIe bridge. This overcomes the limitation on the maximum number of lanes (x8) supported by each Xilinx PCIe core.

The performance measurements show a throughput of 3.4 Gbytes/s for the PCIe Gen2 x8 core with a payload of 256 Bytes. If the dual-core solution is used with two PCIe Gen2 x8 cores, the total throughput reaches 6.9 GBytes/s. The preliminary measurements with the Gen3 single-core had shown a throughput of 6.7 GBytes/s with a payload of 256 Bytes. A custom board based on a Virtex7 FPGA is currently being developed to extend the dual-core architecture to two Gen3 x8 cores.

The DMA engine is currently used in different experimental setups for synchrotron light source applications at ANKA and PETRA III.

Primary author: Mr ROTA, Lorenzo (KIT)

Co-authors: Dr KOPMANN, Andreas (KIT); Prof. WEBER, Marc (KIT); Mr BALZER, Matthias Norbert (KIT); Dr CASELLE, Michele (KIT); Dr CHILINGARYAN, Suren (KIT); Mr STEVANOVIC, Uros (KIT)

Presenter: Dr CASELLE, Michele (KIT)

Session Classification: Poster

Track Classification: Logic