



Contribution ID: 163

Type: Poster

Development of Network Interface Cards for TRIDAQ Systems with the NaNet Framework

Tuesday 27 September 2016 18:20 (1 minute)

NaNet is a framework for the development of FPGA-based PCI Express (PCIe) Network Interface Cards (NICs) with real-time data transport architecture that can be effectively employed in TRIDAQ systems.

Key features of the architecture are the flexibility in the configuration of the number and kind of the I/O channels, the hardware offloading of the network protocols stack, the stream processing and the zero-copy RDMA (for both CPU and GPU) capabilities.

Three NIC designs have been developed with the NaNet framework for the CERN NA62 L0 trigger and for the KM3NeT-IT underwater neutrino telescope DAQ system.

Summary

NaNet is a framework for the development of FPGA-based PCI Express (PCIe) Network Interface Cards (NICs) with real-time data transport architecture that can be effectively employed in TRIDAQ systems.

Key features of the architecture are the flexibility in the configuration of the number and kind of the I/O channels, the hardware offloading of the network protocols stack, the stream processing and the zero-copy RDMA networking capabilities.

Zero-copy RDMA is supported for both CPU and GPU (nVIDIA GPUDirect).

Three different NIC designs have been developed with the NaNet framework for use in the low level trigger of the CERN NA62 experiment (NaNet-1 and NaNet-10) and in the DAQ system of the KM3NeT-IT underwater neutrino telescope (NaNet-3).

Being the most complete of the three in terms of capabilities, we will focus our description on the NaNet-10 design.

Since the beginning of 2016 NaNet-10 has been integrated in the NA62 experiment to implement a GPU processing stage for the real-time generation of refined RICH detectors primitives in order to increase the background trigger rejection and the trigger purity for additional rare decay channels selection.

The target FPGA-based board for this design is the Terasic DE5-NET.

It hosts an Altera StratixV device and allows the integration of up to four 10GbE channels plus a PCIe Gen2/3 x8 host interface.

We implemented the 10GbE channels using both 10GBASE-R and 10GBASE-KR standards to have a wide device compatibility.

Along with the MAC layer we added a complete hardware UDP/IP protocol offloader, enabling minimal latency and full bandwidth for the data channel.

From an architectural point of view, NaNet-10 is a real-time, multiple stream processing system realized through a functional pipeline of hardware blocks executing different tasks on the data streams.

Current partition shows 1) 10GbE interface with UDP transport protocol support (UDP_INTF),

2) highly customizable data manipulation block (STREAM_PROC) and 3) low latency, high throughput RDMA-

based host/GPU interface (NETWORK_INTF).

The UDP_INTF handles the data coming from multiple UDP streams (four in the NA62 RICH detector) and multiplex them in a single channel while the NETWORK_INTF integrates a low latency GPUDirect RDMA hardware engine able to directly inject data in CPU and/or GPU memory.

The STREAM_PROC block is currently customized for the NA62 RICH detector data protocol and executes a sequence of different tasks on the multiple data streams received through the UDP_INTF, eliminating the need of time-costly data re-ordering in GPU memory.

In particular its first stage performs a decompression on the stream data coming from the RICH read-out channels; its second stage reformats the decompressed data to get a “GPU-friendly” alignment of data structures in memory; its last stage executes a time alignment of the different data streams to a common timestamp and produces a single, merged, GPU-aligned stream of the data coming from the entire RICH detector.

Authors: LONARDO, Alessandro (Universita e INFN, Roma I (IT)); BIAGIONI, Andrea (Universita e INFN, Roma I (IT)); Mrs LO CICERO, Francesca (INFN Roma); LAMANNA, Gianluca (Istituto Nazionale Fisica Nucleare Frascati (IT)); PONTISSO, Luca (Universita di Pisa & INFN (IT)); MARTINELLI, Michele (INFN); Mr FREZZA, Ottorino (INFN Roma); VICINI, Piero (INFN Rome Section); AMMENDOLA, Roberto (Universita e INFN Roma Tor Vergata (IT))

Co-authors: ROSSETTI, Davide (Universita e INFN, Roma I (IT)); Mrs PASTORELLI, Elena (INFN Roma); Mr SIMULA, Francesco (INFN Roma); SOZZI, Marco (Universita di Pisa & INFN (IT)); Mr CRETARO, Paolo (INFN Roma); Mr PAOLUCCI, Pier Stanislao (INFN Roma)

Presenter: LONARDO, Alessandro (Universita e INFN, Roma I (IT))

Session Classification: POSTER

Track Classification: Trigger