



Contribution ID: 357

Type: **Talk**

Front-End RDMA Over Converged Ethernet, lightweight RoCE endpoint

Wednesday 23 October 2024 13:48 (18 minutes)

In a DAQ system a large fraction of CPU resources is engaged in networking rather than in data processing. The common network stacks that take care of network traffic usually manipulate data through several copies performing expensive operations. Thus, when the CPU is asked to handle networking, the main drawbacks are throughput reduction and latency increase due to the overhead added to the data transmission process. Networking with zero-copy can be achieved by adding a Remote Direct Memory Access (RDMA) layer to the network stack and making dedicated hardware take care of the burden of the stack handling. Considering the ever-growing demand of larger bandwidth for big data systems, many works point in the direction of implementing network stacks on custom hardware. FPGAs are the natural target for reducing time to market and keeping a low entry-barrier. In this work implementation of RDMA directly on the front-end electronics is explored, in this way it is possible to free part of the computing farm's CPU resources. RDMA over Converged Ethernet (RoCE) is the industry-standard Ethernet-based RDMA solution with a multi-vendor ecosystem, making it the natural choice. This work focuses on the hardware implementation of a stripped-down version of RoCEv2 implementing only the transmitter part of the protocol, enabling its deployment in small FPGA such as the rad-hard parts used in the detector front-end. Preliminary results of resource usage, latency and throughput will be shown.

Author: BORTOLATO, Gabriele (Universita e INFN, Padova (IT))

Co-authors: TRIOSSI, Andrea (Universita e INFN, Padova (IT)); BERGNOLI, Antonio (Universita e INFN, Padova (IT)); BORTOLATO, Damiano (INFN - National Institute for Nuclear Physics); MENGONI, Daniele (Universita e INFN, Padova (IT)); MONTECASSIANO, Fabio (Universita e INFN, Padova (IT)); PAZZINI, Jacopo (Universita e INFN, Padova (IT)); ZANETTI, Marco (Universita e INFN, Padova (IT)); MIGLIORINI, Matteo (Universita e INFN, Padova (IT)); VENTURA, Sandro (Universita e INFN, Padova (IT))

Presenter: BORTOLATO, Gabriele (Universita e INFN, Padova (IT))

Session Classification: Parallel (Track 7)

Track Classification: Track 7 - Computing Infrastructure