



Contribution ID: 35

Type: Oral

Front-End RDMA Over Converged Ethernet, real-time firmware simulation

Tuesday 3 October 2023 09:40 (20 minutes)

Several physics experiments are moving towards new acquisition models. In this work implementation of Remote Direct Memory Access (RDMA) directly on the front-end electronics has been explored, in this way is possible to free part of the computing farm's CPU resources. The work also introduces new verification techniques for verifying RDMA over Converged Ethernet (RoCE) firmware block developed at ETH, including real-time firmware simulation using Verilog Simulator. The result is a stripped down firmware version, allowing its implementation on smaller FPGAs, such as rad-hard parts.

Summary (500 words)

Several physics experiments are moving (or are evaluating the possibility to move) towards new acquisition models. The tendency is to leave the hardware trigger system in favour of a complete or partial acquisition of the front-end data paired with a powerful online software event discrimination. Hardware trigger systems usually have to deal with a tight latency budget due to the narrow readout buffering. To reduce the selection inefficiencies resulting from the adoption of not optimal trigger algorithms due to the limited time budget and online computing resources, the main trigger schema is going to be revisited. The traditional first trigger level is going to be replaced by a hardware pre-processing of the data stream followed by a software online selection.

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 RDMA layer to the network stack and making dedicated hardware take care of the burden of the stack handling.

The main goal of the RDMA implementation in the detector front-end electronics is to move up the adoption of clever networking protocols to the data producer. Therefore, the front-end electronics can initiate the RDMA transfer to the computing farm, eliminating the point-to-point connection between the front-end and back-end allowing the freedom of dynamically switching the routing to the computing nodes according to their processing availability. By appropriately choosing the network protocol for RDMA it is also possible to obtain a two-fold benefit. The possibility of adopting commodity hardware makes the DAQ system reduce reliance on custom hardware and it exploits all the advantages of a mature technology. In this way, the DAQ system gains in scalability and easiness of maintainability.

RoCE is the industry-standard Ethernet-based RDMA solution with a multi-vendor ecosystem, making it the natural choice. In this work the implementation and verification of the main firmware blocks for the realisation of the RoCE endpoint have been explored. A real-time firmware simulation of the RoCE network stack has been developed where real network packets are exchanged between free-running Systemverilog code and the host machine via a TUN/TAP device which emulates a connection with a physical device (FPGA). The second part is devoted to show the verification process of the modified RoCE stack using the tools developed so far such as the novel simulation framework. The lightweight RoCE will be a stripped down version of the already verified firmware allowing the deployment on FPGAs with a low resource pool possible target devices

could be rad-hard FPGAs used in front-end detector boards.

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

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

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

Session Classification: Programmable Logic, Design and Verification Tools and Methods

Track Classification: Programmable Logic, Design and Verification Tools and Methods