Contribution ID: **16** Type: **Poster**

First results of an evaluation of 100Gb/s Ethernet as a future HEP readout link

New-generation physics detectors create a need for high-speed, high-flexibility datalinks in the community. Specific interest lies with commercial standards, compatible with off-the-shelf hardware, therefore replacing custom backends.
We present encouraging first results of an effort evaluating 100Gb/s Ethernet for data readout in the context of typical High-Energy Physics detector requirements. Due to asymmetric data-rate requirements in up- and downlinks, the unidirectional operation of Ethernet is examined. The results are verified with synthetic and realistic traffic, using a series of demonstrators. Additionally, to qualify the radiation hardness of 100Gb/s Ethernet, a statistical analysis is performed based on recent radiation test data.

Summary (500 words)

Current front-end datalink systems in High-Energy Physics (HEP) experiments are generally based on custom link protocols that are not compatible with commercial off-the-shelf (COTS) hardware. Instead, purpose-made electronics based on FPGAs are required to acquire the data and provide a translation layer between the detector and the data centre. Currently there is an increasing interest in using commercial protocols (like Ethernet) as detector datalinks for future detectors. This would enable the use of commercial networking equipment to replace custom backend boards, creating a native interface to the data centre. Thus, first-stage data processing can take advantage of modern technologies and high-performance computing. However, Ethernet has not been developed to meet the specific requirements of the HEP community. Therefore, an explorative effort led by the Electronic Systems for Experiments (EP-ESE) group of CERN has been launched in 2023 to evaluate the possible future use of commercial-grade 100Gbps Ethernet (100GbE) as a link protocol, and we will be presenting here our very first results.

Firstly, we examine the capability of using 100GbE as a unidirectional readout link, since the 100GbE standard contains provisions for speed-asymmetric links. The main motivation for such an approach is the low-bandwidth requirements for control data (that are not foreseen to increase significantly in the future). Implementing a 100Gb/s link in this direction would induce unnecessary complexity and restrictions. Thus, for the time being, we decided to focus solely on the readout path. We have conducted tests to verify the unidirectional operation involving 100GbE COTS network interface cards, switches, and a wide range of standard commercial transceiver modules. The unidirectionality was achieved for a peer-to-peer, single-hop connection using raw Ethernet. Additionally, the unidirectionality was demonstrated for multi-hop setups using IPv4 and UDP to provide networking and transport layers. Supplementary saturation tests have shown satisfactory results using synthetic traffic generated by FPGA development kits. To replace the synthetic traffic with a realistic one, we developed a one-to-one protocol translation from 10.24Gb/s Low-power Gigabit Transceiver (LpGBT) uplink to 10Gb/s Ethernet (10GbE) endpoint. The aim is to aggregate the real front-end traffic through multiple such translators inside a 10GbE-to-100GbE COTS switch. To further mimic a realistic deployment of a front-end system directly connected to COTS switches, we report on our work on packaging the translator in an industry-standard small form-factor pluggable (SFP) module. Having such a module could also be proven useful for other purposes, such as the future integration of systems that will not move to new-generation 100GbE front-end ASICs.

In parallel, we show the first results of our studies concerning the performance of such readout link under radiation environment. To do so, we used recent radiation test data collected with a technology demonstrator 100Gb/s ASIC. Using this data set, we performed a statistical radiation qualification of 100GbE. Additionally, we calculated expected error and link failure rates, taking into account forward error correction measures built into the Ethernet standard and the results are very encouraging. A further effort to normalize these rates to LHC-characteristic radiation profiles and fluxes is underway.

Primary author: STUMPERT, Valentin (CERN, KIT - Karlsruhe Institute of Technology (DE))

Co-authors: KLEKOTKO, Adam; PERRO, Alberto (Universite d'Aix-Marseille III (FR)); HERNANDEZ MONTESINOS, Daniel (CERN); MARTINA, Francesco (CERN); VICHOUDIS, Paschalis (CERN); BARON, Sophie (CERN); BIEREIGEL, Stefan (CERN)

Presenter: STUMPERT, Valentin (CERN, KIT - Karlsruhe Institute of Technology (DE))

Track Classification: Optoelectronics and Electrical Data Links