



Contribution ID: 430

Type: Oral presentation to parallel session

NaNet: a low-latency NIC enabling GPU-based, real-time low level trigger systems.

Tuesday, 15 October 2013 13:30 (20 minutes)

The integration of GPUs in trigger and data acquisition systems is currently being investigated in several HEP experiments.

At higher trigger levels, when the efficient many-core parallelization of event reconstruction algorithms is possible, the benefit of reducing significantly the number of the farm computing nodes is evident.

At lower levels, where typically severe real-time constraints are present and custom hardware is used, the advantages of GPUs adoption is less straightforward.

A pilot project within the CERN NA62 experiment is investigating the usage of GPUs in the central Level 0 trigger processor, exploiting their computing power to implement efficient, high throughput event selection algorithms while retaining the real-time requisites of the system. One of the project preliminary results was that data transfer over GbE links from readout boards to GPU memories using commodity NICs and vanilla software stack consumed the biggest part of the time budget and was the main source of fluctuations in the global system response time.

In order to reduce data transfer latency and its fluctuations we envisioned the usage of the GPUDirect RDMA technology, injecting readout data directly from the NIC into the GPU memories without any intermediate buffering, and the offloading of the network stack protocol management from the CPU, eliminating OS contribution to latency and jitter.

We implemented these two features in the NaNet FPGA-based NIC: the first was inherited from the APENet+3D NIC development, while the second was realized integrating an Open IP provided by the FPGA vendor.

We will provide a deep description of the NaNet architecture and a detailed performance analysis of the integrated system on the NA62 RICH detector GPU-based L0 trigger processor case study, along with an insight of future developments.

Primary authors: LONARDO, Alessandro (INFN, Roma I (IT)); BIAGIONI, Andrea (Universita e INFN, Roma I (IT)); ROSSETTI, Davide (U); Dr LOCICERO, Francesca (INFN); Dr SIMULA, Francesco (INFN); Dr TOSORATTO, Laura (INFN); Dr FREZZA, Ottorino (INFN); Dr PAOLUCCI, Pier Stanislaw (INFN); VICINI, Piero (INFN Rome Section); AMMENDOLA, Roberto (INFN)

Co-authors: PANTALEO, Felice (CERN - University of Pisa); LAMANNA, Gianluca (CERN); SOZZI, Marco (Sezione di Pisa (IT)); FANTECHI, Riccardo (Sezione di Pisa (IT))

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

Session Classification: Data Acquisition, Trigger and Controls

Track Classification: Data acquisition, trigger and controls