Contribution ID: **27**   Type: **poster presentation**

Optimizing the transport layer of the ALFA framework for the Intel Xeon Phi co-processor

ALFA is the common framework of the next generation software for ALICE and FAIR high energy physics experiments. It supports both offline and online processing which includes ALICE DAQ/HLT/Offline and the FairRoot project. The framework is designed based on a data-flow model with message-oriented middleware (MOM) serving as a transport layer. By using multiple data-flows concurrently it facilitates parallel processing which maps naturally to emerging multi-core and many-core computing architectures.

With the introduction of the Intel Xeon Phi co-processor in the industry it is interesting to investigate whether it can be used by ALFA to increase the processing efficiency. The co-processor can be used in three main computing modes. These are offload in which portions of code are accelerated on the device, native where the full program is executed on the device and symmetric where complete tasks are executed on the device and the host processor. For acceleration via offloading there are many competitive platforms such as GPUs and FPGAs. Although Xeon Phi can be used as an accelerator, it is particularly interesting to investigate the possibility of utilizing the co-processor in the *symmetric mode* of operation. Since it is x86_64 compatible it is possible to port complete *task processes* to the device and take advantage of the manycore architecture. It is also worth noting that the next generation of the Xeon Phi codenamed Knights Landing (KNL) will be manufactured in a socket variant as well. The research of using the co-processor as an independent node as opposed to just an offload accelerator can serve as a preliminary study for the future KNL servers.

The software components ported to the device will be connected with the rest of the system via the transport layer therefore there is strong motivation to optimize it for Xeon Phi. The metrics against which it is optimized are **throughput**, **latency** and **energy** consumption, with throughput being the primary target. The two core MOM technologies of choice for ALFA are ZeroMQ and NanoMSG. The out-of-the box versions of these libraries use primarily TCP as the transport protocol which is known to provide limited performance on Xeon Phi in terms of data transfer throughput. In this effort these libraries are extended with support for SCIF, the Xeon Phi native transport protocol over PCIe and additionally with the Co-processor Communication Link (CCL), an RDMA technology used for efficient internode communication. By introducing these extensions we will demonstrate improvements in the data transfer performance by collecting performance monitoring results both in isolation with micro-benchmarks and integrated in the ALFA framework with the respective transport layer benchmark. A successful completion of the optimizations of the transport layer will improve the performance of NanoMSG and ZeroMQ on Xeon Phi. This achievement will potentially make this architecture a viable choice for certain use cases for ALFA which will further enrich its heterogeneous computing capabilities.

Primary author: SANTOGIDIS, Aram (CERN)

Presenter: SANTOGIDIS, Aram (CERN)

Track Classification: Track8: Performance increase and optimization exploiting hardware features