



Contribution ID: 415

Type: **Poster**

Data transfer to remote GPUs over high performance networks

Thursday 27 October 2022 16:10 (30 minutes)

In the past years the CMS software framework (CMSSW) has been extended to offload part of the physics reconstruction to NVIDIA GPUs. This can achieve a higher computational efficiency, but it adds extra complexity to the design of dedicated data centres and the use of opportunistic resources, like HPC centres. A possible solution to increase the flexibility of heterogeneous clusters is to offload part of the computations to GPUs installed in external, dedicated nodes.

Our studies on this topic have been able to achieve high-throughput, low-latency data transfers to and from a remote NVIDIA GPU across Mellanox NICs, using the Remote Direct Memory Access (RDMA) technology to access the GPU memory without involving either nodes' operating system.

In this work we present our approach based on the Open MPI framework, and compare the performance of data transfers of local and remote GPUs from different generations, using different communication libraries and network protocols.

Experiment context, if any

CMS

References

Significance

Author: MARAFI, Ali (Kuwait University (KW))**Co-author:** BOCCI, Andrea (CERN)**Presenters:** MARAFI, Ali (Kuwait University (KW)); BOCCI, Andrea (CERN)**Session Classification:** Poster session with coffee break