Contribution ID: **411** Type: **Poster presentation**

Architectural improvements and 28nm FPGA implementation of the APEnet+ 3D Torus network for hybrid HPC systems

Monday, 14 October 2013 15:00 (45 minutes)

Modern Graphics Processing Units (GPUs) are now considered accelerators for general purpose computation. A tight interaction between the GPU and the interconnection network is the strategy to express the full potential on capability computing of a multi-GPU system on large HPC clusters; that is why an efficient and scalable interconnect is a key technology to finally deliver GPUs for scientific HPC.
In this paper we show the latest architectural and performance improvement of the APEnet+ network fabric, a FPGA-based PCIe board with 6 fully bidirectional off-board links with 34 Gbps of raw bandwidth per direction, and X8 Gen2 bandwidth towards the host PC. The board implements a Remote Direct Memory Access (RDMA) protocol that leverages upon peer-to-peer (P2P) capabilities of Fermi and Kepler-class NVIDIA GPUs to obtain real zero-copy, low-latency GPU-to-GPU transfers.
Finally we report on the development activities for 2013 focusing on the adoption of the latest generation 28 nm FPGAs and the preliminary results achieved with synthetic benchmarks exploiting the implementation of state-of-the-art signalling capabilities of PCI-express Gen3 host interface.

Summary

Primary authors: Dr LONARDO, Alessandro (INFN Roma); Dr BIAGIONI, Andrea (INFN Roma); Dr ROSSETTI, Davide (INFN Roma); Dr LO CICERO, Francesca (INFN Roma); Dr SIMULA, Francesco (INFN Roma); Dr TOSORATTO, Laura (INFN Roma); Dr FREZZA, Ottorino (INFN Roma); Dr PAOLUCCI, Pier Stanislao (INFN Roma); Prof. VICINI, Piero (INFN Roma); Dr AMMENDOLA, Roberto (INFN Roma Tor Vergata)

Presenter: Dr AMMENDOLA, Roberto (INFN Roma Tor Vergata)

Session Classification: Poster presentations

Track Classification: Software Engineering, Parallelism & Multi-Core