



Contribution ID: 150

Type: Poster

Architectural Improvements and Technological Enhancements for the APENet+ Interconnect System

Wednesday 24 September 2014 16:46 (1 minute)

In this paper we describe the latest generation of APENet+ network interface card.

This new APENet+ generation delivers a point-to-point, low-latency, 3D-torus NIC integrated in a PCIe Gen3 board based on a state-of-the-art, 28nm Altera Stratix V FPGA.

The NIC features a network architecture designed following the Remote DMA paradigm and tailored to tightly bind the computing power of modern GPUs to the communication fabric.

For the new APENet+ card we show characterizing figures as achieved bandwidths, eye diagrams and BER obtained by exploiting new high performance ALTERA transceivers and PCIe Gen3 compliancy.

Summary

The APENet+ project delivered a point-to-point PCIe Gen2 interconnect adapter to be employed in hybrid x86+GPU computing clusters with a 3D toroidal network mesh. APENet+ high performance and low latency capabilities are the result of a network architecture designed following the Remote DMA paradigm which is tailored to tightly bind the computing power of modern GPUs to the communication fabric. Doing so on the APENet+ board was possible by employing latest standards for the physical interconnect - SFF-8436 with its QSFP+ modules - and exploit them by means of the vast resources provided by a state-of-the-art FPGA platform like the Stratix IV the board is based upon.

Following up with enhancements in these areas - the most significant being an upgraded interconnect standard (SFF-8665) with new zQSFP plugs and the latest generation for the PCIe standard - means employing an evolved FPGA platform able to integrate and exploit these improvements alongside added functionality. This requires a redesign of APENet+ to support the new PCIe Gen3 protocol, faster transceivers and a much larger amount of programmable resources.

The chosen platform for this redesign exploration is the Stratix V FPGA Development Kit. The board is based on an 28nm technology FPGA (5SGXEA7K2F40C2N) and sports a PCIe x8 connector, two HSMC connectors for exporting high bandwidth links over a daughterboard, a QSFP connector with optical cage and an Ethernet PHY 10/100/1000Mbps copper connector.

A PLDA-proprietary component - QuickPCIe Expert - was chosen as a Gen3-compliant core interface towards the PCIe bus. With this core, either an Avalon or an AXI bus can be selected as user interface; the choice of implementing the latter gives us room for preparing the new APENet+ environment to deal with the impact brought by a full-featured ARM processor which is slated to be one of the most significant additions on FPGA platforms of the next generation.

The new PLDA core also allows for greater freedom when choosing among different strategies for the DMA channels management.

In this work we will show how performances vary as a result of different implementations in this regard; moreover, we will show updated, characterizing figures for APENet+ over the new high performance ALTERA transceivers provided by the Stratix V Development Kit as achieved bandwidths, eye diagrams and BER.

Authors: LONARDO, Alessandro (INFN Rome Section); BIAGIONI, Andrea (INFN Rome Section); ROSSETTI, Davide (Nvidia Corporation); PASTORELLI, Elena (INFN Rome Section); LO CICERO, Francesca (INFN Rome Section); SIMULA, Francesco (INFN Rome Section); OTTORINO, Frezza (INFN Rome Section); TOSORATTO, Laura (INFN Rome Section); MARTINELLI, Michele (INFN Rome Section); PAOLUCCI, Pier Stanislao (INFN Rome Section); VICINI, Piero (INFN Rome Section); AMMENDOLA, Roberto (INFN)

Presenter: BIAGIONI, Andrea (INFN Rome Section)

Session Classification: Second Poster Session

Track Classification: Opto