FPGA - Based Compute Nodes for the PANDA Experiment at FAIR

Wolfgang Kühn, Camilla Gilardi, Daniel Kirschner, Johannes Lang, Soeren Lange, Ming Liu, Tiago Perez (JLU, Giessen), Lars Schmitt (GSI, Darmstadt), Dapeng Jin, Lu Li, Zhen'An Liu, Yunpeng Lu, Qiang Wang, Shujun Wei, Hao Xu, Dixin Zhao (IHEP Beijing, Beijing), Krzysztof Korcyl, Jacek Tomasz Otwinowski, Piotr Salabura (Jagiellonian University, Krakow), Igor Konorov, Alexander Mann (TU München, Garching)

***Abstract*—PANDA is a new universal detector for antiproton physics at the HESR facility at FAIR/GSI. The PANDA data acquisition system has to handle interaction rates of the order of 10**7 /s and data rates of several 100 Gb /s. FPGA based compute nodes with multi-Gbit/s bandwidth capability using the ATCA architecture are designed to handle tasks such as event building, feature extraction and high level trigger processing. Each board is equipped with 5 Virtex4 FX60 FPGAs. High bandwidth connectivity is provided by four Gbit Ethernet links and 8 additional optical links connected to RocketIO ports. A single ATCA crate can host up to 14 boards which are interconnected via a full mesh backplane.**

***Index Terms*— Data acquisition, Field programmable gate arrays, Parallel architectures, Triggering**

I. Introduction

PANDA [1] will be one of the major new detector facilities at FAIR facility for Antiproton and Ion research to be constructed at GSI Darmstadt .PANDA (see fig. 1) is a general purpose detector for antiproton physics and will be located in the High Energy Storage Ring (HESR) providing high quality antiproton beams with momenta up to 15 GeV/c for fixed-target experiments.

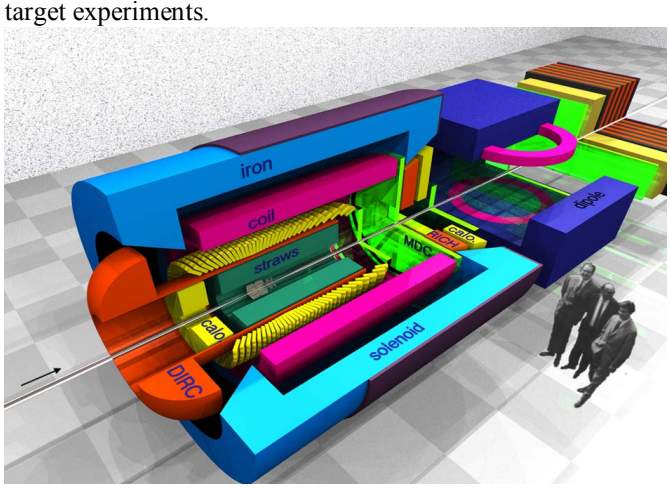

Fig. 1: The PANDA Experiment at FAIR

PANDA features two different spectrometers. The target spectrometer comprises a superconducting solenoid, silicon tracking devices, a central tracker (either TPC or straw tube based), a high quality electromagnetic calorimeter, BaBar-style DIRC for pion/kaon discrimination as well as a muon detector. The forward spectrometer is based on a large normal conducting dipole magnet and drift chambers, Cherenkov detectors and electromagnetic and hadronic calorimeters.

PANDA has a comprehensive physics programme which includes charmonium spectroscopy, the search for QCD exotica, hypernuclear physics as well as the study of charmed hadrons in nuclear matter. This rich and diverse physics programme requires high luminosity experiments with interaction rates in the order of 10^7/s. Raw data rates of the order of 40-200 GB/s need to be processed in real time in order to select events of interest in a potentially very large background. The PANDA Data acquisition system (see fig. 2)

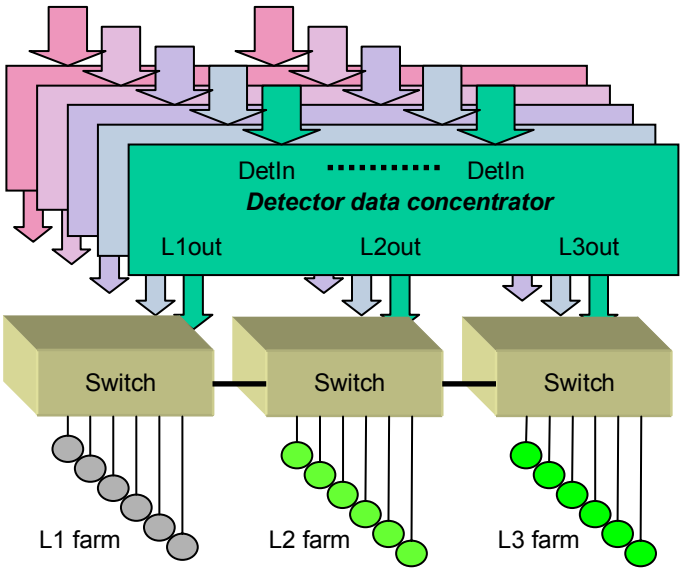

Fig. 2 Schematic overview of the PANDA data acquisition does not employ fixed hardware based triggers but features a continuously sampling system where the various subsystems are synchronized with a precision time stamp distribution

system. Event selection is based on real time feature extraction, filtering and high level correlations. The architecture of the PANDA DAQ system is discussed in more detail in the contribution of Kris Korcyl to this conference.

In the present contribution, we are discussing FPGA based compute nodes (CN) which will serve as a basic building block for the PANDA DAQ system.

II. CN Design Goals and Features

Applications for the CN include pattern recognition for ring imaging Cherenkov detectors, cluster search in calorimeters, tracklet search for the TPC and the silicon vertex detectors as well as more global correlations among subsystems.

FPGA based architectures are most suitable for a cost-effective solution for parallel and pipelined processing. Moreover, modern FPGAs such as the XILINX Virtex 4 FX series provide high speed connectivity via RocketIO as well as via GBit Ethernet.

With such a choice a solution can be build with flexible connectivity both to the detector front-end electronics as well as to conventional LINUX based PC farms for higher level processing.

In order to support algorithms which require for execution more than a single FPGA or even more then a number of FPGAs placed on a single PCB, several CNs could be connected in a way which optimizes the mapping to the data sources. Moreover, the configuration should provide a suitable topology for effective partitioning of the algorithms. The ATCA full mesh backplane is well suited to provide such inter – board connectivity.

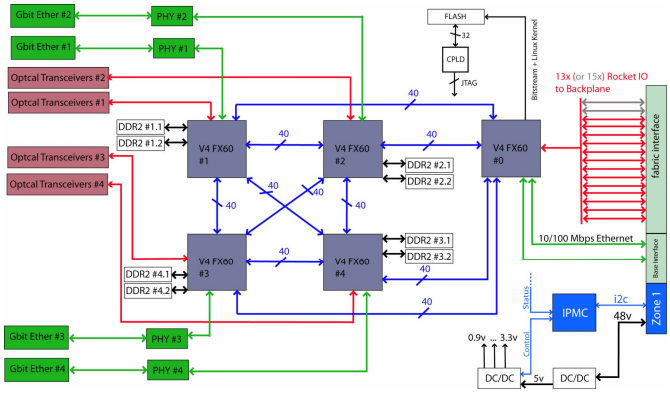

Fig.3: Architecture of the Compute Node

Fig. 3 shows a schematic view of the CN architecture. Its basic building blocks are 5 XILINX Virtex 4 FX60 FPGAs. 4 FPGAs are used to execute the algorithms. The 5th FPGA serves as a switch connecting to other CN modules in the same ATCA shelf via a full mesh backplane.

Each of the 4 processing FPGA is equipped with local DDR2 ram memories for buffering and other purposes such as large lockup tables. Connectivity to the front-end electronics is provided via 4 optical links connected to RocketIO ports on each of the 4 FPGAs. Moreover each FPGA is able to communicate via GBit Ethernet which can be supported either in a standalone mode using the FPGA fabric or using LINUX which is running in the embedded PowerPC processors in the FX60 FPGA.

III. Performance of GBit Ethernet

To get an estimate of the data throughput which could be achieved via Ethernet, we have measured the performance of TCP/IP and UDP transfers using a Virtex 4 FX12 chip on a ML403 board from XILINX. Linux (kernel version 2.6) with a driver supporting the on-chip GBit Ethernet MAC was used. Point-to-point transfers connecting the FPGA and a PC running Linux were performed. We measured transfer rates of up to 300MBits/s for TCP/IP and 375 MBit/s for UDP.

IV. Conclusion

We have presented the design of Compute Nodes which will serve as basic building blocks for the data acquisition and trigger system of the PANDA experiment at FAIR. The most important features include flexible I/O with large bandwidth, providing a scalable solution based on ATCA shelves with full mesh backplane.

The first modules will be available by the end of this year and we plan to set up a small scale system which can be used to upgrade existing experiments such as the HADES experiment at GSI and the BESIII experiment at IHEP Beijing.

Furthermore, small-scale but powerful standalone data acquisition systems could be assembled by combining CNs with a CPU module and ATCA based ADC systems such as discussed in the contribution of Alexander Mann to this conference.

Acknowledgment

This work was supported in part by BMBF and EU.

References

[1] PANDA Letter of Intend, PANDA Technical Design Report (http://www.gsi.de/panda)