A Full Mesh ATCA-based General Purpose Data Processing Board

J. Olsen, T. Liu, Y. Okumura, H. Yin

Fermi National Accelerator Laboratory, P.O. BOX 500, Batavia, IL, 60510, USA

Introduction

High luminosity conditions at the LHC pose many unique challenges for potential silicon based track trigger systems. This is true for both Level-1 and Level-2 trigger applications. Among those challenges is data formatting, where hits and clusters from many thousands of silicon modules must first be shared and organized into overlapping η-ϕ trigger towers due to finite size of the beam's luminous region in *z* and the finite pT curvature of charged particles in the magnetic field. Communication among processing nodes requires high bandwidth, low latency, and flexible real time data sharing. Our hardware design process followed a bottom up approach whereby we studied various track trigger architectures, both existing and proposed. Early in the design process our simulations showed that data sharing between processing nodes is largely asymmetric and highly dependent upon upstream cabling and detector geometry. Implementations involving custom backplanes and discrete cables were considered, however we quickly determined that a full mesh backplane architecture is a natural fit for hardware trigger applications. We have selected the Advanced Telecom Computing Architecture (ATCA) platform as it supports a high bandwidth full-mesh backplane in a robust and reliable industry standard form factor. Our baseline design also works well as a general purpose FPGA-based processor board. The Pulsar II design may prove useful in scalable systems where highly flexible, non-blocking, high bandwidth board to board communication is required.

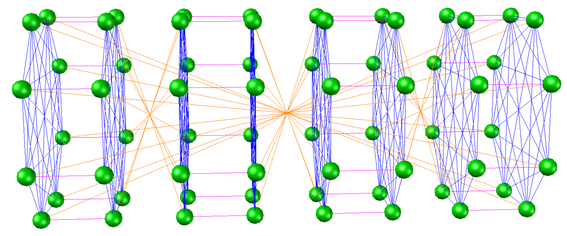

Conceptual view of the ATLAS L2 Fast Tracker (FTK) Data Formatter architecture. The FTK system is organized as 64 trigger towers (4 η x 16 ϕ), represented here by the green nodes. Because the existing silicon tracker was not designed for triggering, the data sharing among trigger towers is complex as indicated by the lines connecting the nodes. This requires the use of a full mesh backplane (shown in blue) for data sharing. Orange lines represent inter-crate links.

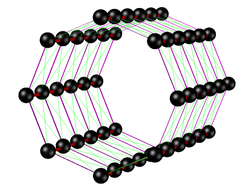

Conceptual view of the proposed CMS phase II L1 tracking trigger towers. The formation is organized as 48 trigger towers (6 η x 8 ϕ). Because the phase II tracker is being designed for tracking trigger purposes, it is possible to arrange the towers in such a way that data sharing only requires communication with immediate neighbor towers. Each node in this diagram represents a trigger tower processor engine. Within each processor engine crate the full mesh backplane is used for time multiplexing of the incoming data, while the simple data sharing is handled with inter-crate links.

AdvancedTCA® Overview

- The **Advanced Telecommunication Computing Architecture** platform has been designed with emphasis on high performance, redundancy and high availability.
- Full Mesh, Replicated Mesh, Dual-Star, and Dual-Dual Star backplanes are available in 6 to 14 slots, rated for 10Gb/s to 40Gb/s per channel.
- Redundant shelf manager boards implement the Intelligent Platform Management Interface. IPMI monitors power and temperature, manages cooling and sequences hot swap insertion and removal.
- Redundant 48VDC power supplies or AC input power.
- Air cooled, up to 500W per slot.
- ATCA front boards are 8U x 280mm and the shelf is 13U in height.

Our 10G full mesh ATCA shelf.

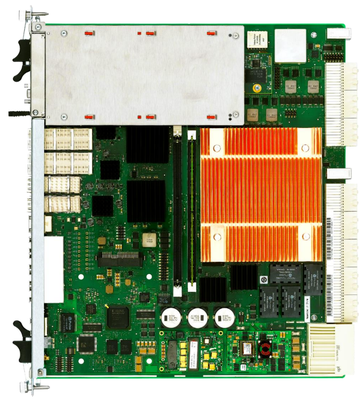

Switch boards reside in slots 1 and 2 and implement a 10G or 40Gb/s Ethernet network. Multiple high bandwidth egress ports are provided on the front panel and RTM. A separate 1Gb/s Ethernet network is also provided for slow controls, downloading firmware, etc. Most switch boards include a multi-core CPU running Linux. Our switch, the Emerson **ATCA-F125** is shown to the left.

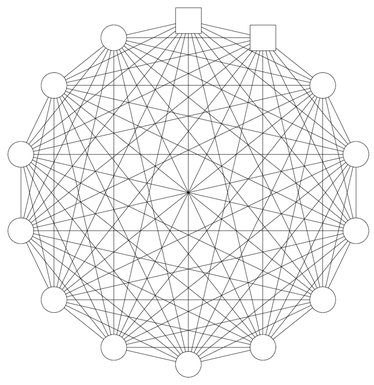

Full Mesh backplanes enable communication between every slot, with no switching or blocking. Each line in this diagram represents a channel which consists of up to four bidirectional ports (lanes).

The ATCA backplane fabric is protocol agnostic. Ports are 100 ohm differential pairs rated for up to 10Gb/s. Custom ATCA boards may use any DC balanced serial protocol such 8B/10B, 1G/10G Ethernet, Infiniband, Fiber Channel, PCI-Express, Aurora, etc.

The Pulsar IIa Prototype

To address the silicon based track trigger needs described above, our design goal is to create a uniquely scalable architecture abundant in flexible, non-blocking, high bandwidth board to board communication channels while keeping the design as simple as possible. Our first prototype, called the Pulsar IIa, is designed around a pair of Field Programmable Gate Arrays (FPGAs). These FPGAs feature multiple high speed serial transceivers which are directly connected to the ATCA full mesh backplane and to pluggable transceivers on the rear transition module (RTM). The Kintex FPGAs we have selected for Pulsar IIa have 16 10Gb/s GTX serial transceivers so our first prototype boards offer a subset of the full backplane and RTM connectivity.

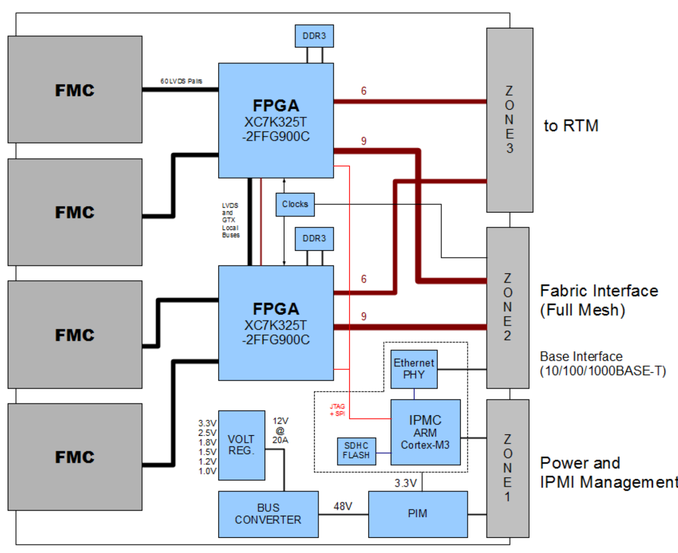

Our FMC Test Mezzanine card features a Xilinx Kintex XC7K160T FPGA, 4 SFP+ transceivers, 128MB DDR3, and a 144 pin socket used for testing custom ASIC chips, such as pattern recognition associative memory devices.

A Cortex-M3 microcontroller is used as an Intelligent Platform Management Controller (IPMC), which is required on all ATCA boards. This microcontroller is responsible for:

- Implementing the IPMI protocol and communicating with the shelf manager board(s), coordinating hot swap operations, etc.
- Running Telnet and FTP servers accessed via the 100BASE-T Ethernet port.
- Managing firmware images on a micro SDHC flash card.
- Programming the FPGAs via JTAG and monitoring over SPI.
- Reading various board temperature and voltage sensors.
- Communicating with the RTM over an I2C bus.

Pulsar IIa Prototype Board

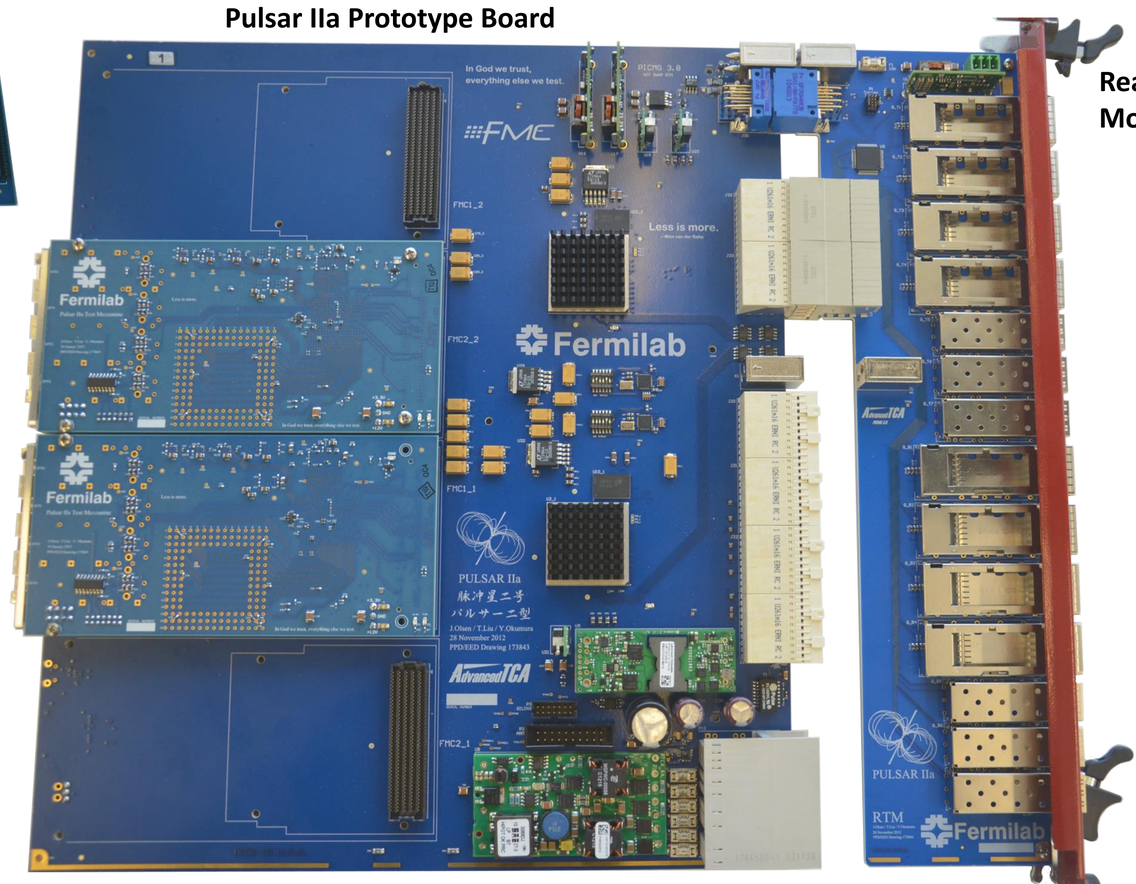

Rear Transition Module (RTM)

Supports up to 8 QSFP+ transceivers. (On the Pulsar IIa each FPGA drives one QSFP+ transceiver.)

Supports up to 6 SFP+ transceivers. (On the Pulsar IIa board each FPGA drives two SFP+ transceivers.)

A small microcontroller on the RTM continuously monitors the state of the transceivers and reports back to the Pulsar IIa microcontroller.

The RTM conforms to the PICMG3.8 standard and supports hot swap.

Pulsar IIa Testing and Results

Our first attempt at designing an ATCA board has proven to be quite successful. The FPGAs are driving data at 10Gb/s through the RTM transceivers and achieving a bit error rates of 10^{-17}. The backplane Fabric Interface channels are stable at 6.25Gb/s (despite the fact that our ATCA backplane is rated for only 3Gb/s).

The Pulsar IIa and RTM boards meet the ATCA mechanical specifications and everything is a good fit in the ATCA shelf. The IPMC microcontroller is communicating with our switch board over the Base Interface Ethernet. The FMC mezzanine cards are transmitting data over LVDS lines to the FPGAs at 400MHz.

ATCA Backplane Fabric at 6.25Gb/s:

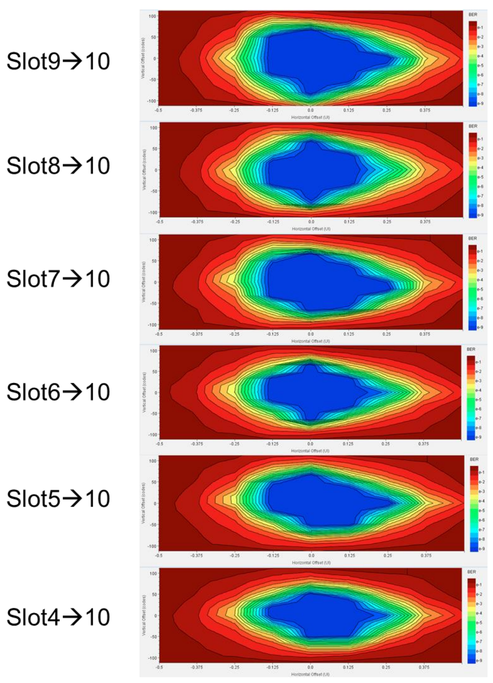

RTM SFP+ Loopback at 6.25Gb/s:

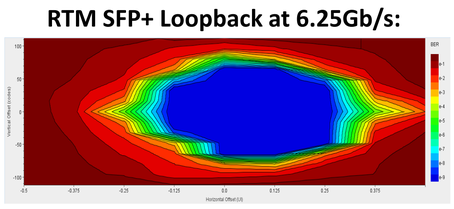

RTM SFP+ Loopback at 10Gb/s:

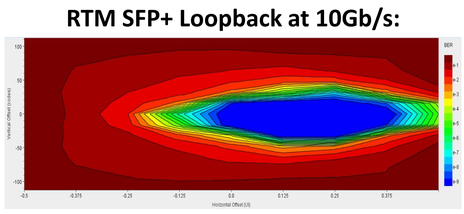

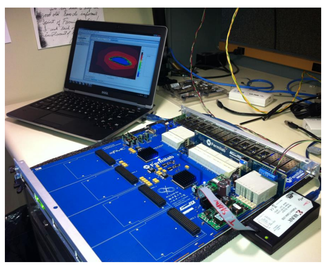

A single slot backplane has been developed for testing boards on the bench top. All Fabric Interface backplane channels loopback and the Base Interface port is brought out to an RJ45 Ethernet jack. A laptop and 48VDC power supply is all that is required to power up and test the Pulsar IIa board and RTM.

Pulsar IIb

Leveraging the experience we gained through designing, building and testing the Pulsar IIa board we are in the final stages of laying out the next generation board, the **Pulsar IIb**. The new board design replaces the two Kintex K325T devices with a single large Virtex-7 FPGA. The GTX transceiver count has increased up to 80 channels, providing a significant bandwidth increase to the RTM, Fabric and Mezzanine cards. The power regulator sections of the board have been redesigned to handle the estimated 30W required by the Virtex-7 FPGA.

The Pulsar IIb design will be used for the ATLAS FTK Data Formatter system. It is also intended to be used for CMS L1 tracking trigger early technical demonstrations.

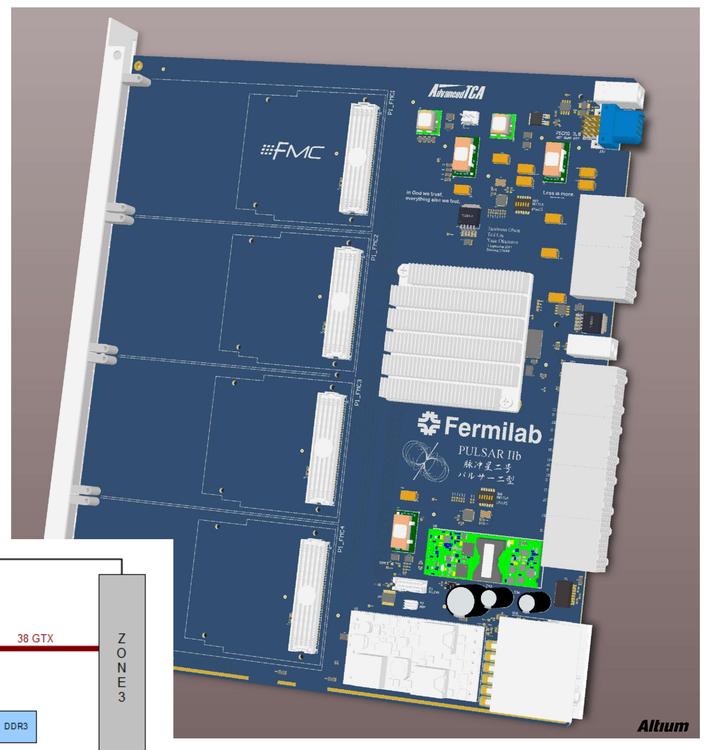

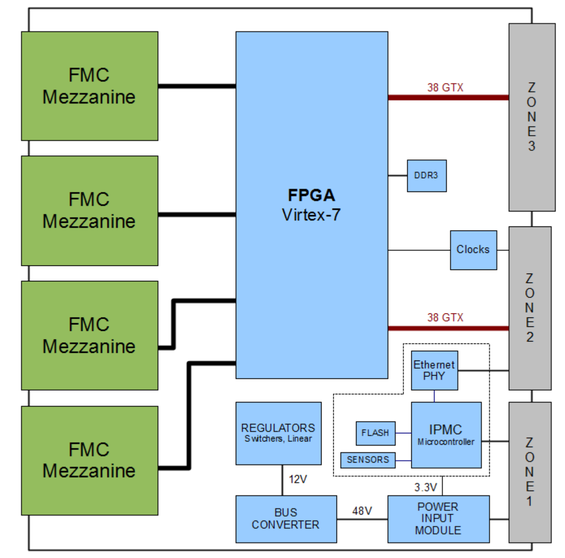

Pulsar IIb Features:

- Designed around a single Virtex-7 FPGA. Supported parts are: XC7VX415T, XC7VX485T, XC7VX550T, XC7VX690T with the FFG1926 footprint.
- Up to 80 10Gb/s transceivers for RTM, Fabric and Mezzanine cards.
- Four FMC Mezzanine card slots; compatible with TTC mezzanine cards from CERN.
- Shelf-wide clock distribution via the backplane Synchronization Interface.
- ARM microcontroller for slow control and IPMC functions.
- 100BASE-T Ethernet Base Interface.
- I2C bus connects various board sensors and SMBus power regulators.
- Up to 512MB DDR3 RAM.
- Hot swap RTM, PICMG 3.8 compliant.
- 14 layer PCB, 8U (322.25mm) x 280mm.

Fermilab