

Development of an ATCA IPMI controller mezzanine board to be used in the ATCA developments for the ATLAS Liquid Argon upgrade

Nicolas Dumont Dayot^a on behalf of the ATLAS Liquid Argon Calorimeter Group

^a LAPP

Annecy Le Vieux, France

E-mail: nicolas.dumont-dayot@lapp.in2p3.fr

ABSTRACT: In the context of the LHC upgrade, we develop a new Read Out Driver (ROD) for the ATLAS Liquid Argon (LAr) community. ATCA and μ TCA (Advanced/Micro Telecom Computing Architecture) is becoming a standard in high energy physics and a strong candidate to be used for boards and crates. We work to master ATCA and to integrate a large number of high speed links (96 links at 8.5 Gbps) on a ROD evaluation ATCA board. A versatile ATCA IPMI controller for ATCA boards which is FPGA Mezzanine Card (FMC) compliant has been developed to control the ROD evaluation board.

KEYWORDS: Large detector systems for particle and astroparticle physics; Data acquisition concepts.



Contents

1. Introduction	1
1.1 Current architecture of the LAr readout	1
1.2 Future architecture of the LAr readout	2
2. Motivations	2
2.1 ATCA platform	2
2.2 ROD evaluation board	3
3. ATCA test board	4
4. ATCA IPMI controller mezzanine	4
4.1 Requirements	4
4.2 Hardware	4
4.2.1 Synoptic	5
4.2.2 Boot	5
4.2.3 EPI interface	5
4.2.4 I2C interface	5
4.2.5 LVDS interface	6
4.3 Software	6
4.3.1 JTAG controller	6
4.3.2 IPMC	6
4.3.3 FRU generator	6
5. Conclusion	7

1. Introduction

In the context of the High Luminosity LHC (HL-LHC) the ATLAS Liquid Argon (LAr) electronics architecture will be redesigned. These modifications will touch the front end electronics (plugged directly on the detector) in addition to the backend electronics which is installed in the counting room.

1.1 Current architecture of the LAr readout

In the current architecture (figure 1), one front end board (FEB) collects signals from 128 calorimeter cells; the detector is equipped with 1524 FEB to read out ~ 180000 channels. Each FEB amplifies using three gains ($\sim 1/10/100$), filters the signal and then samples it at 40MHz. Sampled data are stored in an analog pipeline waiting for the Level 1 (L1) trigger decision during at least $2\mu\text{s}$. The selected data are digitized with a 12 bits ADC (keeping the gain information on 2 extra bits), serialized and sent through one optical link running at 1.6Gbps to the backend electronics.

At this level, the Read Out Driver (ROD) performs E (energy), $E\tau$ (peak timing) and χ^2 (factor form) processing. Each ROD is able to read eight FEB; 192 RODs are connected to the FEB. The results are then distributed to the Read Out Subsystem (ROS) PC.

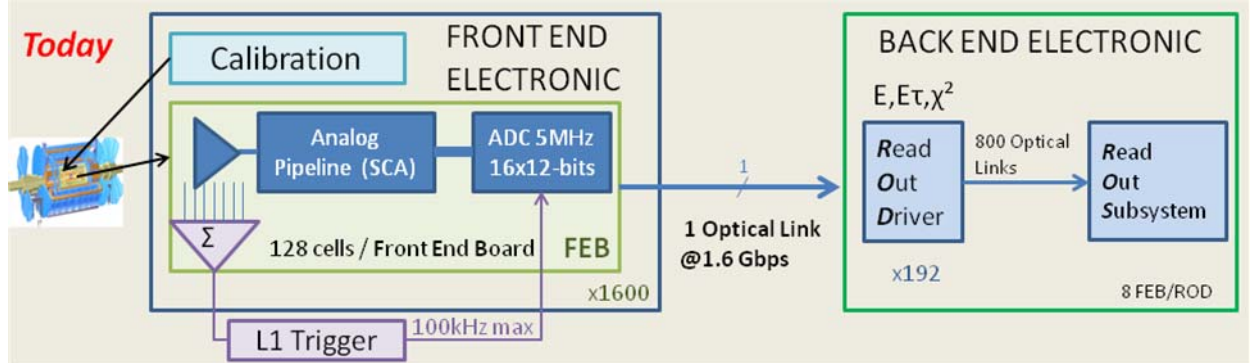


Figure 1 : The current architecture of the LAr readout

1.2 Future architecture of the LAr readout

The current foreseen design is to treat data in the FEB at 40MHz. All the data will be digitized with 18 bits ADC (real 16 bits ADC with auto calibration) and sent to the ROD. The output data rate will increase to reach 100Gbps per FEB. This output bandwidth requires the implementation of 12 parallel optical links running at 10Gbps.

At the backend electronics side, E , $E\tau$ and χ^2 will always be processed but the ROD could also be part of the L0 or L1 trigger decision. This additional feature will imply that the ROD module should be able to accept data from 14 FEB in order to build a complete trigger partition.

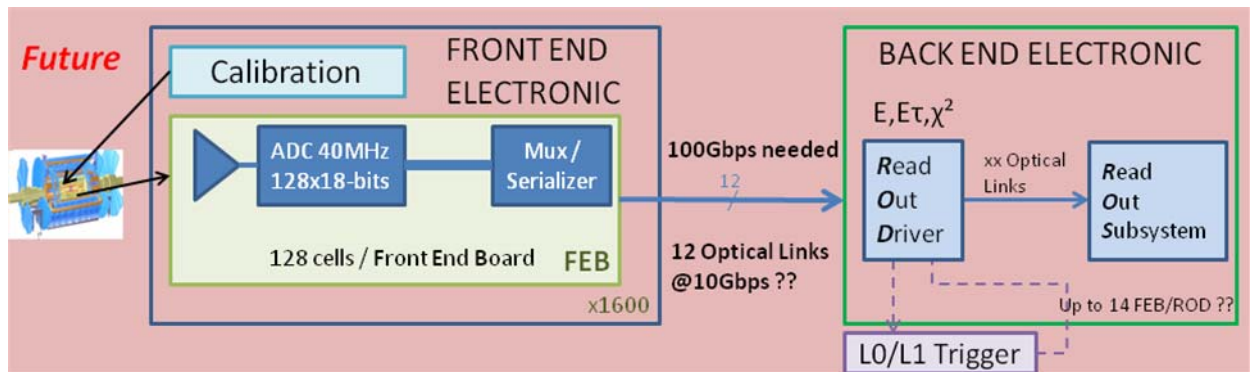


Figure 2 : The future architecture of the LAr readout

2. Motivations

2.1 ATCA platform

For a future upgrade we will need high input and output bandwidth, high speed communications between boards and powerful signal processing at the ROD level. This is why ATCA platform was chosen. This platform provides high density communications between boards over backplane and allows plugging Rear Transition Module (RTM) and Advanced Mezzanine Card (AMC) on a large mother board (32x28 cm). In addition such a standard is reliable with hot swap capability and real time diagnostic.

2.2 ROD evaluation board

We plan to build an ATCA board to evaluate ATCA specifications with Intelligent Platform Management Interface (IPMI) facilities and management through Ethernet. The ROD evaluation board (Figure 3) is a high speed and high density board which integrates the last generation of 12x10Gbps optical fibers and FPGA with high speed transceivers and power DSP cells.

It includes 3 FPGA building blocks which are designed around an ALTERA STRATIX IV GX FPGA implementing 48 transceivers going up to 8.5Gbps. The boot is done with a CPLD reading a parallel 512 Mbits Flash at power up. The FPGA can access to the Flash but also to a 1Gbits DDR3 memory. FPGA 1&0 absorb incoming data from two 12x10Gbps optical receivers and can generate output data with two 12x10Gbps optical emitters to emulate the receivers. So the whole board is able to absorb a data rate of 400Gbps. Eight high speed serial links connect the FPGA0&1 to the FPGA2. The FPGA2 is linked to the ATCA backplane with 4 links on the update interface in order to communicate with it direct neighbour and with 12 links on the fabric interface.

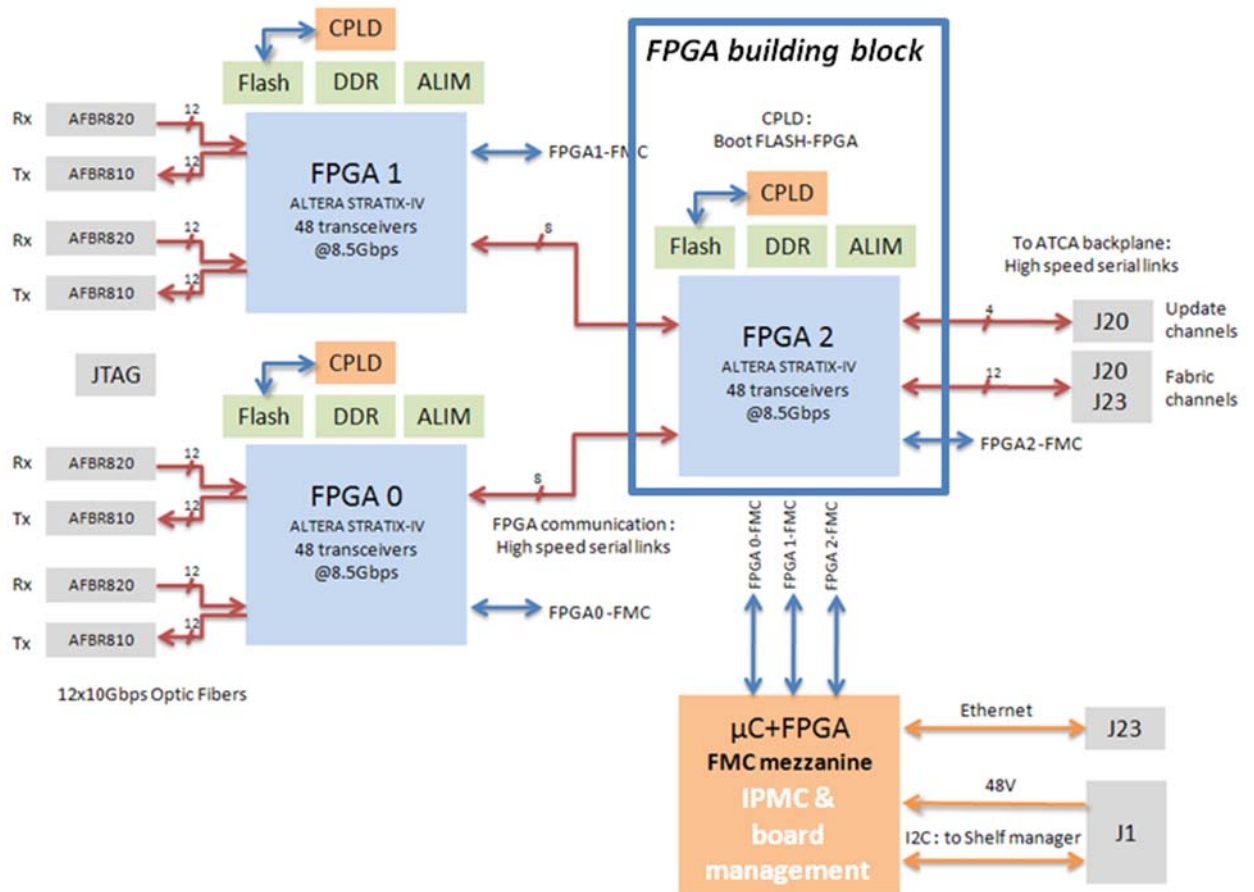


Figure 3: The ROD evaluation board

For the board management, we have designed a FPGA Mezzanine Card (FMC) called ATCA IPMI controller that converses with the external world through Ethernet on the base interface. So we are able to upload the FPGA firmware and other parameters through the network. It can also acts as IPM Controller (IPMC) to satisfy the ATCA shelf rules and standard. This is described in detail further.

3. ATCA test board

Before launching the production of the ROD evaluation board we want to test the ATCA IPMI controller mezzanine and the FPGA building block on a mother board called ATCA test board.

The core of the FPGA building block (figure 4) is a low cost ALTERA ARRIAI GX FPGA which incorporates the same functionalities than the STRATIX IV GX FPGA. The building block has been fully tested. So the boot from the CPLD and the parallel Flash works and allows us to store several firmwares in this memory. The FPGA can write in the FLASH and can access the DDR3 memory with a data rate up to 1Gbps.

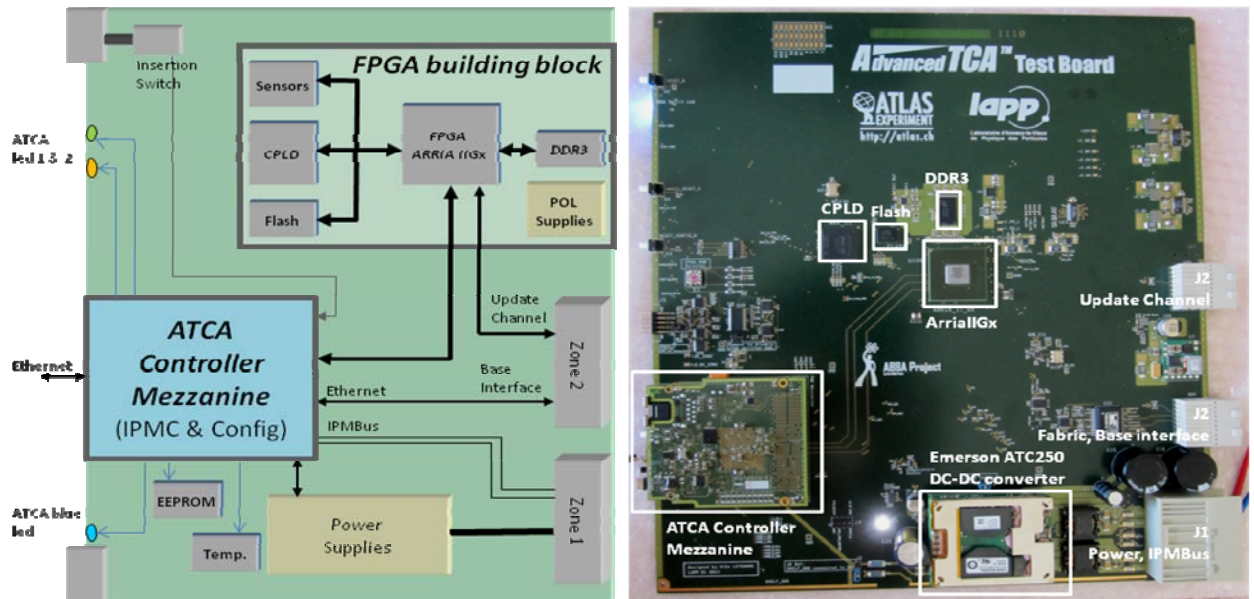


Figure 4: The ATCA test board

4. ATCA IPMI controller mezzanine

4.1 Requirements

One of the requirements of this project is to manage the board through Ethernet on the base interface (J2 connector): we want to upload the FPGA firmware or other parameters (optimal filtering coefficients) and to monitor sensors, registers or memories through the network. We also want to program the FPGA with an embedded JTAG controller driven by Ethernet. The dialog between the mezzanine and the FPGA is mainly done with serial LVDS links.

This mezzanine acts also as IPM Controller through 2 redundant I2C IPM Buses in order to communicate with the shelf manager which handles all the boards in the ATCA crate. As it is drawn on figure 4, the IPMC is principally interfaced with an insertion switch for the hot swap detection; with an EEPROM storing board information; with the enable pin of the DC/DC converter which provides the 12V; with an ATCA blue led indicating the state of the board and with a temperature sensor. It can also monitor the FPGA voltage or temperature sensors through the LVDS links. This will be useful for alarm or failure diagnostic.

4.2 Hardware

This board fulfills the FPGA Mezzanine Card (FMC) standard format with high pin count (400 pins): 160 links with 74 differential links are available on this type of connector.

4.2.1 Synoptic

This board is mainly designed around one microcontroller and one FPGA. The power supply is the IPMI 3.3V given by the ATCA DC/DC converter on the mother board.

The FPGA is a SPARTAN6 XILINX FPGA with highly configurable IO. All the FMC inputs/outputs are driven by the FPGA in order to be configured by the users. The FPGA boots from a SPI FLASH and exchanges data with the microcontroller through a parallel bus called External Peripheral Interface (EPI bus : 16 bits data associated with 12 bits address).




Figure 5: ATCA IPMI controller mezzanine

So the FPGA is the door opened toward the mother board: VHDL designers can implement dedicated protocol (I2C, SPI, LVDS links ...) with the external world for their own specifications.

The microcontroller is an ARM cortex M3 processor from Texas Instrument (LM3S9B92) and has multiple interfaces. JTAG and USB interfaces help the developers to debug. The Ethernet link comes from either the front panel RJ45 connector or from the base interface. This processor includes also several functions like EPI bus management, IPMC, JTAG controller....

4.2.2 Boot

The FPGA boots correctly from the SPI FLASH which can be uploaded directly by the microcontroller. USB/JTAG and Ethernet have been validated: the microcontroller code and the FPGA firmware can be uploaded through the external network with the RJ45 connector but not through the base interface for the moment.

4.2.3 EPI interface

The read/write accesses to/from the FPGA register/memories works well with the EPI communication. The parallel bus cycles are cadenced by a 25MHz clock.

4.2.4 I2C interface

In addition an I2C VHDL master reads the temperature sensor each second and stores its value in a register in order to be monitored by the IPMC module. The EEPROM I2C interface is still to be developed.

4.2.5 LVDS interface

The LVDS links have been tested up to 400Mbps with the mother board FPGA through the FMC connector. A protocol between the FMC and the external FPGA still needs to be thought about (8B/10B code,header,data,checksum....)

4.3 Software

LINUX is the operating system used for the software development: GNU GCC compiler, GNU GDB or OpenOCD debugger and Texas Instrument programmer are the main tools to design code for the microcontroller and its different modules or libraries. So JTAG, IPMC, IP stack and hardware libraries have been designed. In addition a friendly user environment will provide a web interface, a TCP/IP client and a file server.

4.3.1 JTAG controller

The JTAG chain can be mastered either by external JTAG connector or by the microcontroller configured as JTAG controller. In that case some GPIO are used for TDI, TDO, TMS and TCK emulation. So it provides an embedded JTAG controller driven by Ethernet. An .xsvf files player has been validated to program the FPGA whereas .jam files have been abandoned to program the FLASH through the CPLD: the microcontroller does not have sufficient memory space for .jam files and the programming time is too long (~45").

The JTAG controller is also able to configure the JTAG chain multiplexer (SCANSTA111 from National Semiconductor) located on the mother board. This allows separating the JTAG chains when there are multiple FPGA, CPLD on the mother board.

4.3.2 IPMC

The IPMC fullfills the following specifications : PICMG 3.0R3.0 ATCA base specifications and IPMI v1.5 with relevant subset of IPMI v2.0. The code is based on the open source CoreIPM and suits only for an ATCA board not a mother board with AMC: for the moment it does not support an IPM Local (IPML) bus for the AMC management.

The ATCA power up state machine has been validated. This state machine must never fail and must handle properly the insertion switch, the hardware address, the ATCA blue led, the DC/DC enable, the temperature sensor and the 2 IPM buses linked to the shelf manager. This code has been under stress test with power up/down cycles every 5" in order to detect reliability issues. The procedure has been stopped without any failure after 48 hours.

Now the next step is to design an event generator in order to send information to the shelf manager when alarms or failures occur on the board (voltage limits, high temperature...).

4.3.3 FRU generator

Field Replaceable Unit (FRU) data collect information related to the identity of the board like the IO description on the J2 connector: for example, number and location of these channels with their protocol are detailed. The Sensor Data Record (SDR) describes the sensor of the board like the insertion switch, the temperature sensor....This information is stored in the mother board EEPROM in order to be read by the IPMC and sent to the shelf manager at power up.

An FRU generator (figure 6) has been developed in order to build the entire C structure of the board. This generator processes the definition area which is kind of dictionary of the main definition and the user definition with a M4 pre-processor. Then it builds the whole C structure of the board. This tool is a generic tool and can be used for all ATCA board users.




Figure 6: FRU generator

Then this C structure will be the starting point for binary the file uploaded into the mother board EEPROM. Next step will be to develop the same tool for the SDR generator.

5. Conclusion

We have designed an ATCA ROD board to evaluate the functionalities needed by the next generation of the ROD for the ATLAS LAr upgrade. This board is a high speed and high density board with large number of fast links and recent generation of FPGA and optical receivers on an ATCA board.

For the control of that board we have designed a generic ATCA IPMI controller mezzanine in the FMC format. It can act as IMP controller and has direct Ethernet access for uploading FPGA firmware and other parameters and monitoring registers and memories through the network.

This mezzanine and its software are currently tested on an ATCA test board we have designed. Tests performed have been successful, but much work still needs to be done.