

Next generation ATCA control infrastructure for the CMS Phase-2 upgrades

Aleš Svetek, Jes Tikalsky, Marcelo Vicente¹, Robert Fobes, Sridhara Dasu, Tom Gorski and Wesley Smith

University of Wisconsin-Madison

1150 University Ave., Madison WI 53706-1390, USA

E-mail: marcelo.vicente@cern.ch

A next generation control infrastructure to be used in Advanced TCA (ATCA) blades at CMS experiment is being designed and tested. Several ATCA systems are being prepared for the High-Luminosity LHC (HL-LHC) and will be installed at CMS during technical stops. The next generation control infrastructure will provide all the necessary hardware, firmware and software required in these systems, decreasing development time and increasing flexibility. The complete infrastructure includes an Intelligent Platform Management Controller (IPMC), a Module Management Controller (MMC) and an Embedded Linux Mezzanine (ELM) processing card.

*Topical Workshop on Electronics for Particle Physics
11 - 14 September 2017
Santa Cruz, California*

¹Speaker

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

<https://pos.sissa.it/>

1. Introduction

Several institutes are presently designing hardware, firmware and software that will be used in CMS for data taking during the High Luminosity LHC (HL-LHC) physics runs, starting in 2025 [1]. Most of the architectural choices have been focused in designing Advanced TCA (ATCA) blades and, as part of the PICMG 3.x standard [2], each blade needs to implement a set of control and management functionalities.

The next generation control infrastructure presented includes three different hardware systems that together implement the PICMG 3.x standard, provide additional features, flexibility and decrease the time necessary to design and test an ATCA blade.

1.1 Embedded Linux Mezzanine

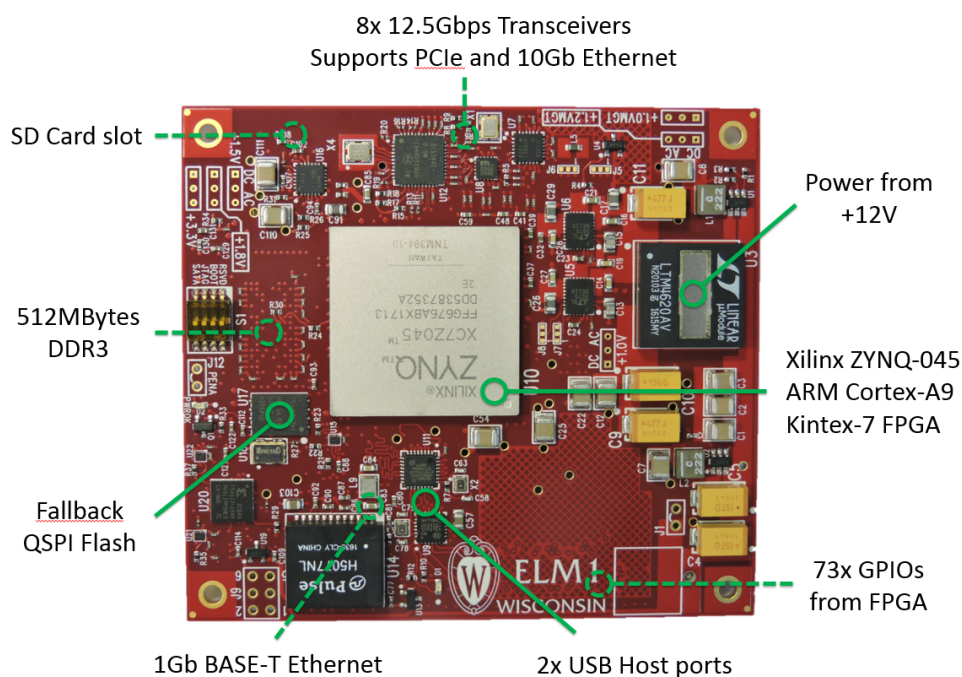


Figure 1: Embedded Linux Mezzanine (ELM)

The Embedded Linux Mezzanine (ELM), depicted in Figure 1, is intended to be used as the blade's on-board computer and its primary access point. It features a dual core ARM processor packed inside a high-end Xilinx ZYNQ System-on-Chip (SoC) running a full-fledged Linux distribution. Expandability is provided by the ZYNQ's FPGA section where 73 general purpose input/outputs (GPIOs) and 8 multi-gigabit transceivers (MGTs) are available to the hardware designer and can drive several standards and custom protocols, such as:

- JTAG Master Controller (Xilinx Virtual Cable),
- AXI Chip2Chip (parallel high-speed GPIOs, Aurora),
- PCI Express (x1, x2, x4 and x8 lanes),
- 10Gbit Ethernet (XFI, XAUI).

Blade peripherals are connected to GPIOs or MGTs and configured using Xilinx ecosystem, predominantly made up of Xilinx Vivado and standard AXI inter-connectivities. Common peripherals can be other FPGAs (e.g. Xilinx Ultrascale+) or simpler devices such as EEPROMs, allowing them to be accessed and configured through Linux or exposed by Ethernet. Figure 2 exemplifies how the ELM can be connected to other peripherals.

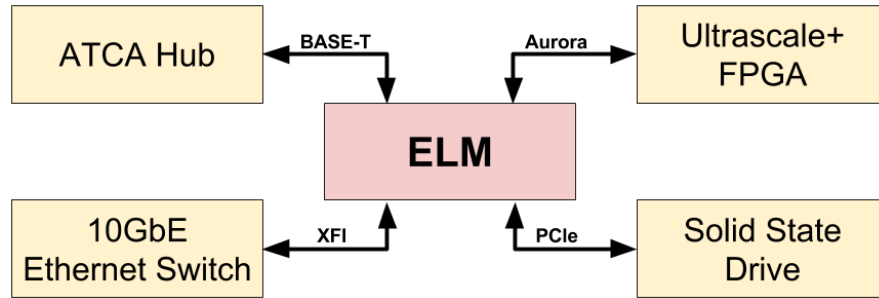


Figure 2: ELM connectivity example

1.1.1 Motivation for ELM development

Although there are existing embedded solutions such as COMe [3] mezzanines equipped with x86 Intel processors, the ELM provides additional flexibility with its ZYNQ SoC which fully integrates FPGA logic and embedded processing at multiple levels, including the physical package, the firmware and the device driver level. As an example, the same boot process that loads the CPU operating system also programs the FPGA fabric.

The ELM form factor and FPGA logic can also serve as a customizable 10GbE endpoint tying into the main board, which the COMe cannot easily provide given the age of the platform and the lack of customizable inputs/outputs.

Lastly, the pinout of COMe is set by PICMG, and the options available target support of the PC/AT IO connectivity, much of which is not needed on processing boards in HEP applications. The ELM retains use of the 220 pin COMe connector, and some IO pin definitions such as USB, console and 1GbE, but is able to retask other pins from superfluous IO functions (extra USBs or LVDS display ports) to programmable IO pins that can be used to support the FPGA(s) and associated peripherals on the main board in HEP applications.

1.2 Intelligent Platform Management Controller




Figure 3: Intelligent Platform Management Controller (UW-IPMC)

The Intelligent Platform Management Controller (UW-IPMC) mezzanine, shown in Figure 3, is intended to be the next generation of blade controllers that communicate with ATCA Shelf Managers (ShMC). It is part of the PICMG 3.x specification and each blade needs to have one of such controllers. They are responsible for doing blade related health checks (e.g. temperature, voltages) and to act in case of problems, such as over-temperature or over-voltage.

The IPMC design has a low-cost Xilinx ZYNQ SoC that runs FreeRTOS, which is ideal for time sensitive applications, and uses FPGA logic to do parallel comparisons of sensor data versus threshold values, allowing the IPMC to trigger on alarms at sub-millisecond rates. The firmware, depicted in Figure 4, uses FreeRTOS queues coupled with a publisher/subscriber broker system to relay important messages across different tasks, such system allows for easy debugging and logging when needed.

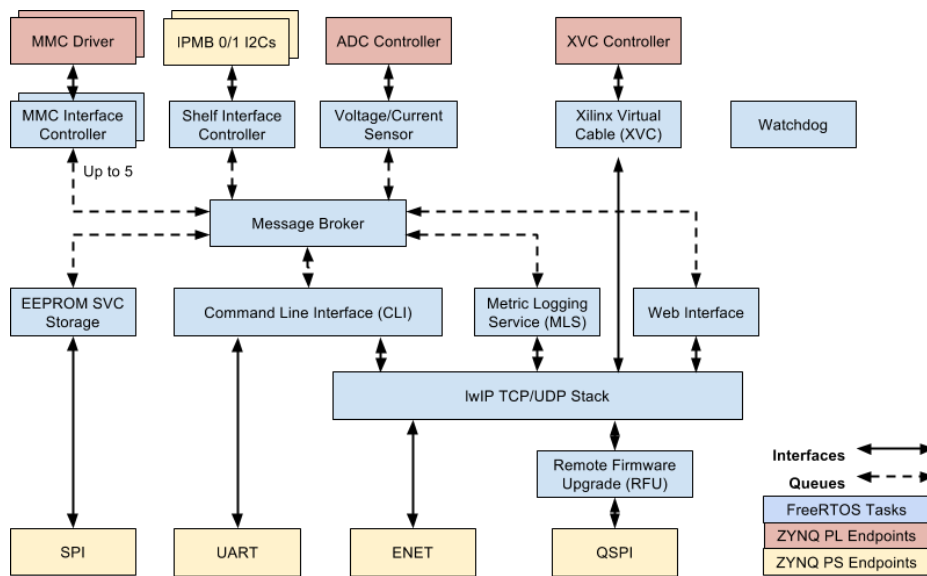


Figure 4: UW-IPMC firmware and software architecture

Five independent Modular Management Controller (MMC) interfaces allows the IPMC to manage Advanced Mezzanine Cards (AMCs), this prevents bus locking in case one AMC fails during operation. For expandability, 49 GPIOs are available at the mezzanine connector and wired to the FPGA section, allowing the IPMC to interface with power, monitoring or management related components on the ATCA blade. GPIOs can also be used to add more MMC interfaces as required. Access to the IPMC is done through Ethernet via the ATCA Hub card or through the IPMI bus.

1.2.1 Motivation for UW-IPMC development

The UW-IPMC project is primarily motivated by three desired features not available in the LAPP/CERN IPMC [4] pinout as projected for HEP applications:

First, the LAPP/CERN pinout allocates support for 9 subsidiary controllers (MMCs) on either an ATCA Rear Transition Module or front board Advanced Mezzanine Cards (AMCs), a

number seems larger than necessary. To require 9 controllers, each AMC would have to be of the smallest size possible (half-height, single width in PICMG terminology), and cards of that size are seldom seen in High Energy Physics as new designs. The UW-IPMC allocates pin space to support 5 controllers, freeing those pins on the connector for other functions.

Second, all 9 MMCs in the LAPP/CERN pinout share a single communications bus, a multi-master I2C connection. In the event that one of the MMCs develops a problem, it could affect operations for all modules on the bus without giving a clear indication as to which is the failing unit. The UW-IPMC, using the programmable logic of the ZYNQ SoC device to create the necessary I2C peripherals, uses the same approach as commercial MicroTCA MCH modules, and dedicates a separate physical IPMB connection to each MMC.

Lastly, the LAPP/CERN IPMC allocates a single I2C connection to be used by all main board peripherals to implement sensors. This leads to a comparatively slow and crowded communications bus on which it is not possible to poll individual sensors much faster than a few times per second, if the sensor count is high. Thus if non-recoverable fault conditions occur on the board (say an overvoltage), the IPMC is not in a position to quickly intervene, possibly preventing permanent damage to the main board. The UW-IPMC has 16 channels of on-board ADC capability at 1kSample/sec or faster, which can detect faults quickly.

1.3 Module Management Controller

The next generation Module Management Controller (MMC) is provided as a reference design and not as a mezzanine due to its simplicity. The design moves away from RISC architectures to ARM microprocessors by using Atmel's SAM4N. Similarly to IPMCs, several ADCs and GPIOs allow the MMC to manage and monitor AMCs which are pluggable and hot-replaceable modules in ATCA blades. Rear Transition Modules (RTMs) are also supported.

The firmware implementation uses FreeRTOS for time sensitive tasks, such as sensor monitoring and alarming. Communication with the IPMC is done through a dedicated I2C bus.

References

- [1] Contardo, D., Klute, M., Mans, J., Silvestris, L., Butler, J., *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, HCC-P-008
- [2] PICMG, *AdvancedTCA Base Specification 3.0*, Revision 3.0
- [3] PICMG, *COM Express Module Base Specification*, Revision 3.0
- [4] J. Mendez et al., CERN-IPMC Solution for AdvancedTCA Blades, these proceedings, 2017