ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA071

PLANS AT CERN FOR ELECTRONICS AND COMMUNICATION IN THE DISTRIBUTED I/O TIER

Grzegorz Daniluk*, Evangelia Gousiou,
CERN, Geneva, Switzerland

Abstract

Controls and data acquisition in accelerators often involve some kind of computing platform (VME, PICMG 1.3, MTCA.4...) connected to Distributed I/O Tier electronics using a fieldbus or another kind of serial link. At CERN, we have started a project to rationalize this tier, providing a modular centrally-supported platform which allows equipment groups to focus on solving their particular problems while benefiting from a set of well-debugged building blocks. The paper describes the strategy, based on 3U Euro crates with a generic FPGA-based board featuring space for FMC mezzanines. Different mezzanines allow communication using different protocols. There are two variants of the electronics, to deploy in environments with and without radiation tolerance requirements. The plans we present are the result of extensive discussion at CERN among all stakeholders. We present them here with the aim of gathering further feedback and potential interest for inter-lab collaborations.

INTRODUCTION

The control system for CERN's accelerator facility is built of multiple layers of hardware and software. These tiers are spanning from the hardware deployed close to the machine, up to the software running on computers that operators use for control and monitoring. As Figure 1 shows, one can distinguish the following three hardware layers:

- Front-end Tier - PLC or a powerful computer in various form factors (VME, PICMG 1.3, MTCA.4, etc.) running an operating system and a set of user applications controlling Distributed I/O Tier electronics over a fieldbus.
- Fieldbus Tier - a networking solution that ensures communication between the master in the Front-end tier and a set of slaves in the Distributed I/O Tier
- Distributed I/O Tier - electronics modules installed close to the machine in radiation-exposed or radiation-free areas controlled by the master in the Front-end tier over the fieldbus. These are usually FPGA-based boards sampling digital and analog inputs, driving outputs and performing various safety critical operations.

Depending on the needs of a given application, one can use either off the shelf equipment (like PLCs with remote I/O modules), design completely custom electronics or have a mix of the two. The very particular needs of the CERN accelerators often demand the development of custom electronics. For these, our group already provides a centralized

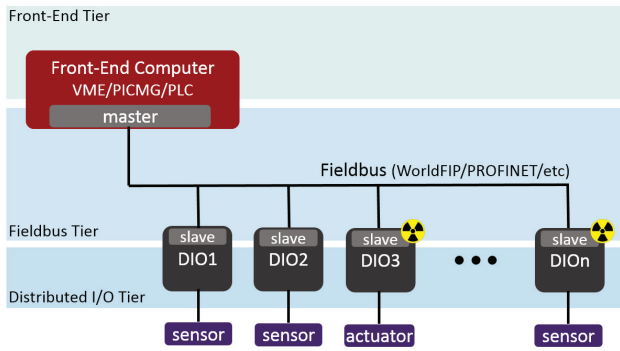

Figure 1: Three hardware tiers of the CERN's control system.

service in the Front-end Tier. It comes in the form of VME crates and PICMG1.3 computers that can host our modular FMC (FPGA Mezzanine Card) kit [1]. However, we currently miss a centrally supported service for modular and reusable electronics in the Distributed I/O Tier. Each equipment group so far (Power Converters, Cryogenics, Machine Protection, Beam Instrumentation, and others) has independently developed custom solutions for their needs. While those devices are different for each application, many of the needs are the same and we have identified that they could also be dealt with centrally.

In the Fieldbus Tier our current services for custom electronics are mainly built around the radiation-tolerant WorldFIP bus. With this project we also want to enhance our offering, to include the Ethernet-based fieldbuses widely used in industrial applications.

Our project does not aim at replacing all the already installed equipment, but rather targets new developments and foreseen upgrades.

INDUSTRIAL FIELDBUSES FOR ACCELERATOR CONTROL SYSTEMS

Ethernet's steady march over the decades has also found its way into the industrial environment as a means to interconnect field devices to the first level of automation. Protocols such as PROFINET, EtherNet/IP, POWERLINK, EtherCAT usually build on top of standard Ethernet so as to ensure deterministic data exchange and synchronization over the network.

Industrial Ethernet accounts for 46% of the market share of fieldbuses, according to 2017 research from HMS Industrial Networks [2]. The classic non-Ethernet fieldbus position remains strong at 48% of the market with

* grzegorz.daniluk@cern.ch

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA071

an annual growth rate of 7%, thanks largely to Profibus' dominance with 17% market share. Ethernet however is growing faster than previous years, with a growth rate of 22%. PROFINET is first and runners-up are EtherCAT and POWERLINK. Various gateways exist that provide interoperability between the different protocols. CERN follows this global pattern and today the majority of fieldbus installations in radiation-free zones are Profibus while a few newer installations are on PROFINET. For areas closer to the LHC machine, the radiation-tolerant WorldFIP bus is used.

Regarding radiation-free areas, after having evaluated the different Industrial Ethernet protocols we concluded that all of them could technically fit the most demanding future use case, as described in Table 1. In order to provide flexibility to the users and not restrict the protocol range, it was decided to build hardware that supports the four main protocols through one single chip like the NetX52 [3]. With such a chip the hardware remains the same for different protocols. Areas subject to radiation have the same requirements in terms of fieldbus use and the radiation tolerance levels described in Table 2.

Table 1: CERN Fieldbus Requirements

characteristic	**value**
topology	mainly daisy chain
nodes	100
cycle time	1 ms
synchronization	100 us
data-per-cycle	1 Kbyte

For such areas an off-the-shelf chip implementing complex protocols like PROFINET or EtherCAT would not work. Our proposed approach in this case is an in-house implementation of a POWERLINK slave in a radiation-tolerant FPGA. POWERLINK is the only fieldbus with Free and Open Source implementation of the FPGA firmware and software stack [4]. It is the simplest Industrial Ethernet protocol and very much like the WorldFIP, with which we have long experience. Our target FPGA is a ProASIC3 or IGLOO2 from Microsemi. A radiation-tolerant implementation of POWERLINK in a Flash-based Microsemi FPGA has been preferred over the radiation-tolerant-by design fieldbuses (SpaceWire and SpaceFiber) based on the fact that POWERLINK would allow for interoperability with a variety of off-the-shelf Industrial Ethernet nodes in radiation-free zones.

Table 2: Maximum Expected Levels of Radiation in CERN LHC Areas where Electronics are Installed

characteristic	**value**
TID	400 Gy
σ	$<$ 1e-12 cm^2

INTERCHANGEABLE MEZZANINES

In order to facilitate the fieldbus integration to custom electronics in the Distributed IO Tier, while maintaining the principles of modularity as introduced in the Front-end Tier [1], it was decided to implement the fieldbus protocols in a set of FMC mezzanine boards (Fig. 2).

For the radiation-exposed areas the offering includes WorldFIP for simple applications (Fig. 2a) and POWERLINK for those requiring higher bandwidth or interoperability with off-the-shelf PLCs (mock-up in Fig. 2b). In both cases the implementation of the protocol takes place in the radiation-tolerant Microsemi FPGA on the mezzanine. The plan is also to leave part of the FPGA resources unused, so as to allow for user-defined application-specific logic. For the radiation-free areas a single mezzanine housing a multi-protocol chip, like NetX52, could be used for PROFINET, EtherCAT, POWERLINK or EtherNet/IP nodes (mock-up in Fig. 2c). The FPGA in this case is configuring the NetX52 chip leaving resources also to house simple user-logic. The set of mezzanines is completed with the White Rabbit one (Fig. 2d). White Rabbit [5] is a fully deterministic Ethernet-based network for control data transfer and sub-nanosecond synchronization. It is extensively used at CERN and beyond for critical applications demanding very high synchronization accuracy. Note that the White Rabbit protocol operates in the physical and data link layers of the OSI model. It does not define higher layers as the aforementioned Industrial Ethernet protocols.

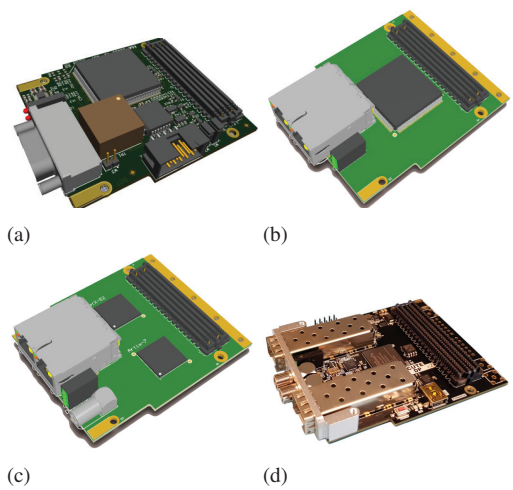

Figure 2: Interchangeable mezzanines for radiation-free and radiation-exposed applications.

DISTRIBUTED I/O TIER MODULAR ELECTRONICS KIT

First Idea - Pizza Boxes

Our first approach for shared electronics in the Distributed I/O Tier was to design a so called "pizza box". It is an already established concept of deploying electronic boards

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA071

enclosed in a 19" rack-mountable box with integrated power supply. Such an enclosure usually has the height of one or two rack units (where 1 rack unit is 44.50 mm). In the frame of the Distributed I/O Tier project we aimed to provide users with two hardware variants: for radiation-exposed and radiation-free areas. Each box would include a generic FPGA-based board in the form factor of the MTCA.4 Advanced Mezzanine Card (AMC) [6]. The AMC has FMC slots and a Rear Transition Module (RTM) connector. One of the FMCs is used to plug a mezzanine for fieldbus communication, as described in the previous section. The other FMC, together with the RTM connector, would be left for the user to customize the box by attaching any application-specific boards.

The idea initially seemed to be very flexible and able to cover various applications, but was finally abandoned. The first issue was the need for more modularity and I/O density per number of 19" rack units. For systems where precise measurements need to be performed, good channel isolation is required and therefore not many analog channels can be squeezed on a single RTM board. One could of course deploy multiple pizza boxes in the same rack, but controlling only a few channels from one FPGA would be a waste of rack space and resources.

The other argument against the pizza boxes is heat dissipation. In the radiation-exposed areas, where access is very restricted and service interventions are limited, it is better to avoid mounting fans for cooling every single box. In that case, stacking multiple horizontally mounted boards would trap warm air between the pizza boxes. This is completely different in a crate setup where boards are installed vertically and warm air is free to flow to the top of a rack.

Modular Crates with a Simple Inter-board Communication

An alternative that addresses the pizza-box limitations, is the modular crate approach. Many of the equipment groups' developments currently deployed in the Distributed I/O Tier are 3U Euro crates (Fig. 3). A Euro crate is a 19" rack

Figure 3: 3U custom Euro crate example.

mountable box, 3 rack units high, that hosts multiple Eurocards connected together through a custom backplane. The standard defines mechanical concepts like the dimensions and card slides in the crate, and leaves open the choice of backplane connectors and the communication bus between the cards. Components to assemble such a chassis are rather cheap and easily available from the local CERN store, which makes it a very attractive solution for various modular electronics applications. The flexibility of Euro crates gave equipment groups over the years a great way of making their custom designs according to the needs of their particular application. However, in the absence of a centralized support in the Distributed I/O Tier, this freedom caused also a proliferation of electronics solutions.

Currently, equipment groups use custom backplane designs differing among the groups and projects. The connector system used to plug the cards is always a multi-row DIN 41612 (the same one used in the VME crates), which is suitable for the low speed needs of the currently deployed systems. Although designing a custom 3U backplane is feasible for slow communication, it becomes challenging to ensure proper signal integrity for high-speed (tens of MHz) transmission. The DIN 41612-type connector itself is also not well suited for fast signaling. The use of differential pairs and single-ended signals with additional ground pins for the return currents could be a solution, but the connector does not provide enough pins to do that.

The target of the new modular electronics in the Distributed I/O Tier is to create a CERN-wide standard for high speed communication over the backplane for these 3U Euro crates. An important principle of the project is to stay close to the standards used in the industry to benefit as much as possible from the work already put into designing the backplanes and simulating their signal integrity. The characteristics of several off-the-shelf backplanes, used also in the Front-end Tier of various control systems, were thoroughly examined:

- **MicroTCA.4** was excluded based on the mechanical requirements. Distributed I/O Tier electronics are often part of safety-critical systems installed close to a very expensive accelerator equipment, like bending magnets or power converters. Dry contact or short circuit in the backplane connector may prevent a system from detecting a quenching magnet, therefore destroying the magnet. MicroTCA uses an edge connector which densely packs 170 pads in two rows. It requires precise PCB manufacturing and careful card insertion to guarantee good contact. The other two standards (PXIe, Compact PCI Serial) use a more reliable multi-row pin connectors which we prefer for this kind of applications.
- **PXIe** defines more reliable multi-row pin connectors, but has an active backplane. For radiation-free areas this would not have been an issue, but in this project we need to cover also installations in the radiation-exposed zones. For these, the best solution is a completely passive backplane - containing only differential and single ended lines. Otherwise, any active component in the

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA071

backplane would have needed to be qualified under radiation and could be a potential source of unreliability.

- **Compact PCI Serial** defines an industrial off-the-shelf fully passive backplane. It is used in various transportation applications [7]. It has robust mechanics to mount cards in the chassis and plug them to a backplane using AirMax VS, a multi-row pin connector. The connector has proven its reliability in various vibration and shock tests[1].

All the backplanes listed above utilize modern communication protocols like PCI-Express, SATA, USB and Ethernet. These are however overcomplicated for the Distributed I/O Tier electronics as the peripheral cards can be very simple devices. In particular, they can just contain opto-isolators for digital I/O lines, or ADCs/DACs for analog channels.

The plan for the common Distributed I/O Tier modular electronics is to reuse the Compact PCI Serial backplane as a set of differential and single-ended lines. It could be used either in the existing 3U Euro crates or with an off-the-shelf chassis, depending on the project needs. The fully passive backplane enables using it with a simple (and better suited for Distributed I/O Tier applications) communication bus. Fast SPI (Serial Peripheral Interface bus) channels are well suitable for the main, point-to-point data exchange between the boards (Fig. 4). It is a full-duplex bus which is simple to implement in an FPGA (also for the radiation-tolerant variant) and can work even with hundreds-MHz clock. It is also a solution widely used in the current CERN custom electronics. SPI gives the flexibility of plugging both the simple (e.g. including only ADCs, DACs or GPIO Expanders) and more complex boards (e.g. including FPGA for data preprocessing). The former would have the ADC/DAC chips connected directly to the SPI interface in the backplane. The latter could use the SPI to receive commands from the System Board depending on the desired application. These fast SPI buses are specified, to ensure interoperability between Peripheral Boards designed by different groups. One could for example foresee that a radiation tolerant ADC or GPIO Expander board designed for one application could be useful for many other Distributed I/O Tier systems.

If, for a given application, the SPI interface is not a good choice, the user can also implement any other communication bus according to the specific needs, because the System Board will have the backplane lines connected to the on-board FPGA.

Besides the star topology of the SPI channels going from the System to Peripheral Boards, a shared Service Bus and multi-drop interlock (interrupt) lines are also foreseen - see figure 4. The Service Bus would contain a slow shared SPI bus (to expose various management and diagnostics information like the board ID or the temperatures readout) and a JTAG chain to be able to remotely reconfigure FPGAs on Peripheral Boards.

[1] Available on request from the connector and backplane manufacturers.

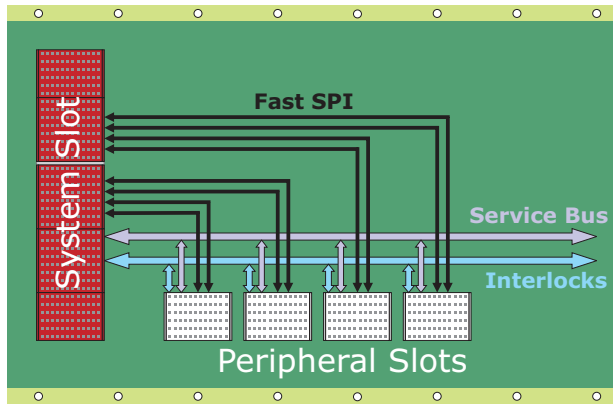

Figure 4: Distributed I/O Tier backplane example (4-ports).

The boards hierarchy would remain according to the Compact PCI Serial standard i.e. one main System Board in a crate communicating with the fieldbus and exchanging data over the backplane with multiple Peripheral Boards. Figure 4 presents a 4-port example (for simplicity) of such a backplane.

The project includes the design of two System Boards:

- for radiation-free areas - System-on-Chip (SoC) based, using a powerful Xilinx Zynq SoC or Xilinx Zynq UltraScale+ MPSoC with an FPGA and a multi-core ARM for data collection and processing according to a user application, with DDR memory and FMC slot
- for radiation-exposed areas - Flash FPGA based, using a Microsemi ProASIC3 or IGLOO2 FPGA for data collection and processing according to a user application, with FMC slot.

Both boards need to have an FMC slot to enable fieldbus communication using one of the previously described FMC mezzanines. The processing power of these two System Boards will be very different. That is because also the needs in radiation-free areas are much different from these for the radiation tolerant electronics. For the safety-critical systems running in the radiation-exposed zones the complexity of electronics and data processing is reduced and pushed to the Front-end Tier. Actions performed by these systems can be described in four simple, but critical steps: sample inputs, check if the values are above some predefined threshold, drive outputs, send postmortem data over a fieldbus to the Front-end Tier. On the other hand, much more computationally expensive applications can take place in the radiation-free areas, like for example the power converters regulation.

ISBN: 978-3-95450-193-9 doi:10.18429/JACoW-ICALEPCS2017-THPHA071

THREE SERVICES FOR HARDWARE IN DISTRIBUTED I/O TIER

We plan to gather the hardware solutions described in the previous sections into three centrally supported services for the Distributed I/O Tier custom electronics:

- Communication chips
- Interchangeable mezzanines
- 3U crate with generic FPGA-based boards

The most sharing and reuse between the equipment groups can be achieved using a 3U crate with a generic FPGA-based board. A new development could be based on this form factor and that would reduce the work flow to a design of custom Peripheral Boards specific to a given application, and focus on firmware implementation for the provided FPGA-based main board.

For systems that cannot benefit from this new form factor, a custom FMC carrier in any desired format can be designed to use one of our FMC mezzanines. This way the system can still benefit from a centrally supported fieldbus communication. All the mezzanines will have FPGA resources available to the final user, therefore they can be used as the main system FPGA. This will reduce the PCB design effort to a simple, standalone, application-specific FMC carrier.

Finally, if none of the above solutions can be used, we also plan to provide a CERN-wide support for the communication chips (nanoFIP [8] for WorldFIP and NetX52 [3] for Ethernet-based fieldbuses). In that case, developers can benefit from our reference designs, HDL IP-Cores to configure and interface with these chips, and all the expertise on the fieldbus communication.

CONCLUSIONS

At CERN we are moving towards a centralized support for custom electronics in the Fieldbus and Distributed I/O Tier of the control system. We plan to develop a modular service consisting of a set of interchangeable FMC mezzanines and a 3U crate with two generic FPGA-based boards to cover the needs in radiation-exposed and radiation-free zones. By providing different building blocks we target the fast integration of the modules by different users that either develop new systems or renovate existing ones. With this project we also bring industrial standards to the world of custom Distributed I/O Tier electronics at CERN. The interchangeable FMC mezzanines implement Industrial Ethernet fieldbuses (PROFINET, EtherCAT, EtherNet/IP, POWERLINK) for radiation-free areas, complemented by WorldFIP and POWERLINK for radiation-exposed areas. The modular crate proposal improves a well established 3U Euro crate standard by using an off-the-shelf Compact PCI Serial backplane for high speed inter-board communication. Finally, two FPGA-based 3U System Boards (for radiation-free and radiation-exposed areas) hosted in the Euro crates provide a generic platform for application-specific HDL modules designed by the equipment groups.

We believe that the modular approach described in this paper can be applied also in other control systems. We want to trigger the interest outside CERN for providing a centralized support for custom electronics in the Fieldbus and Distributed I/O Tier and build a new inter-lab collaboration.

REFERENCES

[1] E. V. der Bij, M. Cattin, E. Gousiou, J. Serrano, T. Włostowski, "CERN's FMC kit", *Proceeding of ICALEPCS WECOCB01, San Francisco, CA, USA, 2013.*

[2] Industrial network market shares 2017 according to HMS, https://www.hms-networks.com/press/2017/02/20

[3] Hilscher NetX52 chip, https://www.hilscher.com/products/product-groups/network-controller/asics/netx-52/

[4] POWERLINK source code, http://openpowerlink.sourceforge.net

[5] White Rabbit project page, https://www.ohwr.org/projects/white-rabbit

[6] Advanced Mezzanine Card Base Specification, *PICMG Std. AMC.0 R2.0.*

[7] MEN, Nürnberg, Germany Compact PCI Serial products and applications, https://www.men.de

[8] nanoFIP project page, https://www.ohwr.org/projects/nanofip