A MODULAR APPROACH TO ACQUISITION SYSTEMS FOR FUTURE CERN BEAM INSTRUMENTATION DEVELOPMENTS

A. Boccardi, M. Barros Marin, T. E. Levens, B. Szuk, W. Vigano, C. Zamantzas
CERN, Geneva, Switzerland

Abstract

This paper will present the new modular architecture adopted as a baseline by the CERN Beam Instrumentation Group for its future acquisition system developments. The main blocks of this architecture are: radiation tolerant digital front-ends, a latency deterministic multi gigabit optical link, a high pin count FMC carrier used as a VME-based back-end for data concentration and processing. Details will be given on the design criteria for each of these modules as well as examples of their use in systems currently being developed at CERN.

THE CERN BEAM INSTRUMENTATION GROUP

The Beam Instrumentation (BI) Group is responsible for designing, building and maintaining the instruments that allow observation of the particle beams and the measurement of related parameters for all CERN accelerators and transfer lines, comprising ultra-relativistic and non machines, circular and linear accelerators and even an antiproton decelerator. Each machine has very specific characteristics and requirements and therefore rarely the same instrument can be deployed on several of them. Typical beam parameters to be measured and monitored are: position, loss, intensity (total and per bunch), tune and chromaticity, luminosity, transfer lines matching, transverse and longitudinal profiles. The BI group decided to adopt a modular architecture for its acquisition system to promote the reuse of electronics and code, to simplify the maintenance and the design of this large variety of instruments.

Despite the details of each implementation can be very different the generic architecture of an instrumentation acquisition system is in the majority of the cases similar (see figure 1): the signal from the sensor is first conditioned and then digitized before being in-line processed.

Figure 1: Typical structure of an instrumentation acquisition system. The signal from the sensor is conditioned, digitized and in line processed.

For the majority of BI systems the processing is performed in a back-end unit, installed in a radiation safe area, even for the cases in which the digitalization is performed at front-end level. This allows the standardization of the back-end itself. That is one of the main block of this modular architecture, and reduce the complexity of the rad-tolerant electronics eventually required, with a consequent increase of reliability.

A problem common to all the system with a remote digitalization is the transmission of the data from the front-end to the back-end with a constant delay to avoid the need of continuous recalibrations at each reset of the system. This is particularly difficult for systems requiring large bandwidth and therefore complex serialization protocols relying on synchronization and PLLs. BI decided to adopt the multi-gigabit and latency deterministic optical link designed in the CERN PH-ESE group as the second main block of its new design toolkit, the third one being a digital front-end module implementing this link in a rad-hard version for all the systems where the digitalization stage is not only remote but also exposed to radiations.

Those three elements, on which the system specific blocks plugs, are described in the next chapters.

THE BACK-END

The board designed to be the standard back-end for the BI applications is the VFC-HD [1]. It is an FPGA-based 6U VME 64x module with a high pin count FMC (VITA 57) slot, the possibility to connect to custom VME rear transition module (RTM), 6 small form factor pluggable (SFP) slots, 4 of which dedicated to user applications and 2 to system ones, and on board DDR3.

The board was specified to be an FMC compliant carrier for its users to be able to use commercially available mezzanines. The decision to adopt the high pin count standard is linked to the need to use the latest generation of fast ADCs and DACs. Nowadays those adopt multi gigabit serial communication for their interfacing and therefore requiring the use of lines not defined in the low pin count standard of the connector.

The RTM was specified to suit the needs of systems requiring part of their connectivity to be plugged in the back of the crate.

The decision for the VME standard came from the need of being able to integrate the new board with existing systems. BI made in the past years extensive use of custom VME crates for its developments. In those crates the beam synchronous clocks and triggers are distributed by a timing receiver over the back plane and the P0
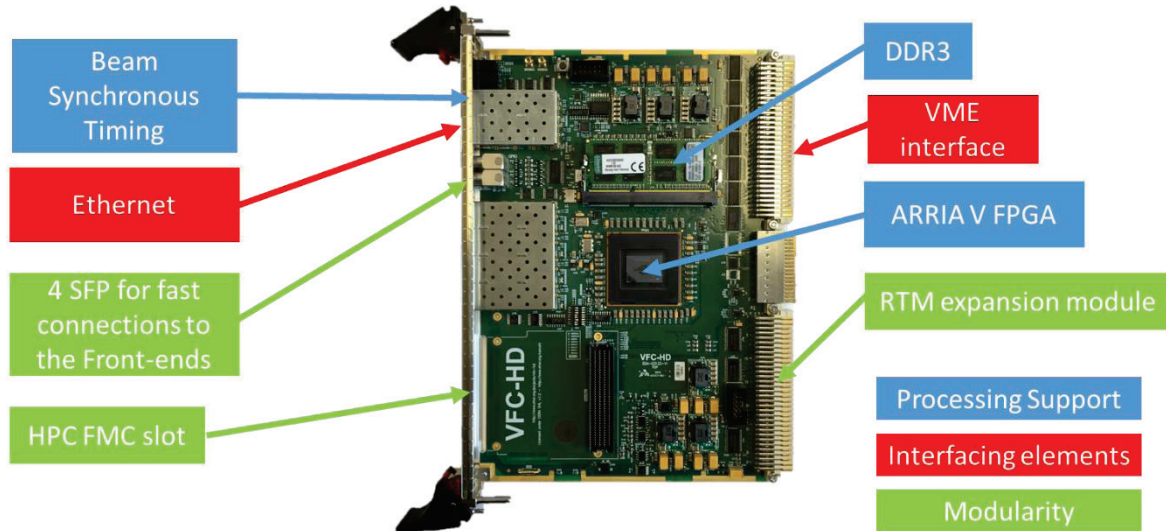
Figure 2: The VFC.

connector to the other VME modules. This led also to the decision to keep the BI specific definition of the P0 signalling in for the VFC-HD.

New systems on the other hand will be based on standard crates, where no timing distribution is implemented on the back plane. Each VFC-HD will need to receive and decode by itself the beam synchronous timing via an optical link, the same way was done by the timing receiver for the whole crate. The first of the system SFP connections is dedicated to this purpose.

In view of possible changes in the timing distribution protocol the VFC-HD was designed with White Rabbit [2] node capabilities, as this is one of the most likely to be new timing distribution protocol.

The VFC-HD can also work in standalone mode to suit highly distributed systems like for example the tune or the instability trigger ones. The board for this purpose can be accessed via an Ethernet connection implemented using the second of the system SFPs.

The 4 application SFPs are meant to connect the back-end to remote digital front-ends like the GEFE(see next chapters). Those SFPs are connected to the FPGA multi-gigabit transceivers in a scheme that takes into account the clocking requirements for the implementation of the chosen latency deterministic link (see next chapter). A single VFC-HD can connect to up to 4 of those front-ends without the need of a FMC optical connection module.

The decision to use on-board DDR3 memories was not immediate. Indeed this is the second high pin count carrier developed in BI: the first one used SRAM chips. The need for high density memories comes from few systems requiring long histories of raw measurements for post mortem analysis or specific beam studies. This requirement came after , i.e. the FPGA model was already chosen i.e. the ARRIA V GX. Neither this model nor the pin compatible GT version supports the fly-by topology used in the DDR3 DIMMs, but to limit the redesign effort

and the cost, an important factor for the major users, was decided not to change it. The total memory on board is 512MB, considered enough, and could be further extended to 2GB if the compatibility with the new "TwinDie" chips from Micron will be confirmed.

THE LATENCY DETERMINISTIC LINK

The latency deterministic link chosen by BI for the communication between its high bandwidth frontends and the VFC-HD is based on the GigaBit Transceiver (GBT) link [3], developed by the CERN PH-ESE group. For non-radioactive environments, this link can be implemented in FPGAs in conjunction with commercial optical modules. For systems that are exposed to radiation, the GBT link is implemented using dedicated rad-tolerant chipsets, comprising an encoder-serializer/decoder-deserializer chip, the GBTx, and a custom optical module, the VTRx.

The latency determinism was a key parameter in the choice of this link. It allows to look at the link as a transparent pipe for the data with a constant delay that can be calibrated just once at the commissioning of the instrument instead of at each reset of either of the two communication ends.

Another element that made of this link the chosen one in BI is the reliability, indeed it features a protocol (GBT Frame) with a robust error detection and correction algorithm (Reed-Solomon) without the need of further interventions of the users.

Additionally, the GBT link also offers high payload bandwidth, 3.2Gbps for a 4.8Gbps communication, the availability of the radiation tolerant chipset, and finally the possibility to use as base clock the beam synchronous 40MHz of the LHC or SPS, making it available as recovered one on the front-end side.
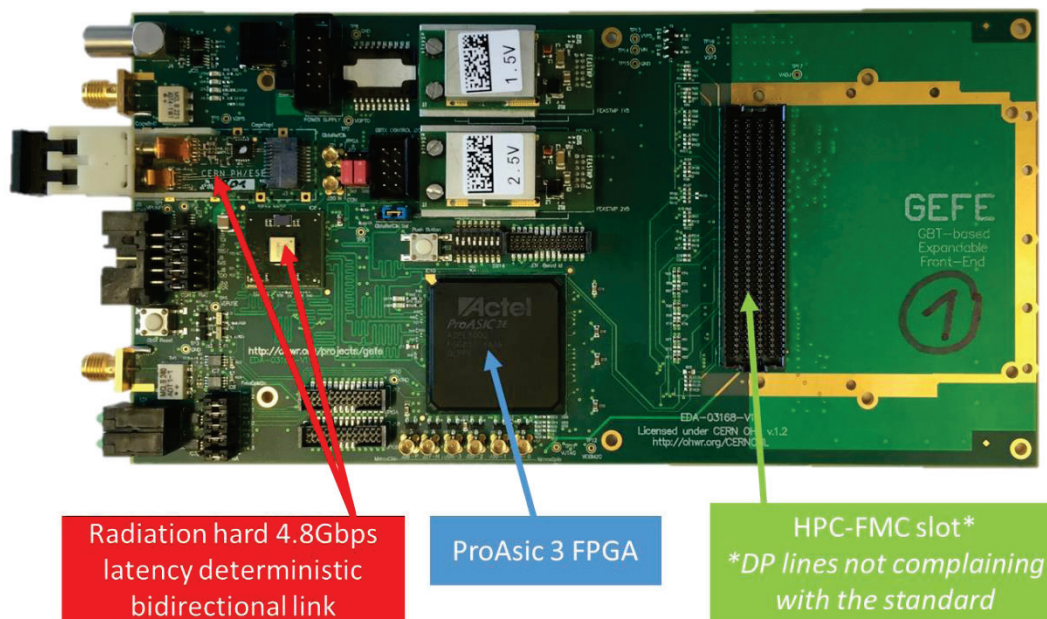
Figure 3: The GEFE.

THE RADIATION TOLERANT DIGITAL FRONT-END

The GBT-based Expandable Front-End (GEFE) [4] is a multipurpose FPGA-based radiation tolerant board designed by BI to be the new standard FMC carrier for radiation tolerant digital front-end applications. Its intended use ranges from fast data acquisition systems to slow control installed close to the beamlines. In the case of the GEFE the choice of the FMC format is for consistency with the VFC-HD and not in view of the use of commercial mezzanines.

It is important to note that the GEFE's FMC connector is for high pin count mezzanines, but that the multi gigabit lines are not implemented in accordance with the standard.

The main purpose of the GEFE is the implementation of the interface between the back-end and the digitalization module.

All the components chosen for this board are either radiation tolerant by design, like the GBTx chipset [3] and the power supplies [5] designed in the PH-ESE group, or have been qualified for a total ionising dose (TID) up to 75krad, a limit imposed by the ProAsic3 FPGA from Microsemi.

The presence of an FPGA limits the TID but increases the flexibility of the board. It allows an easy implementation of the glue logic required to interface the rest of the front-end system to the GBTx chip, to receive and decode commands and configurations from the back-end and to manage the clocking scheme.

The last important factor that was taken into account during the design phase was the size of the board. The aim was to have a module as compact as possible to ease

the integration with other modules close to the beam, where often the space is limited. The GEFE is a dense 200mm by 100mm board.

SYSTEM EXAMPLES

This section introduces examples of systems already using or planning to use the proposed architecture.

The LHC DC Beam Current Transformer (BCT) and the AWAKE (a new experimental line for wake field acceleration) Beam Position Monitoring (BPM) systems both use the VFC-HD and have their sampling in a remote location not exposed to radiation. Both implement their communication between the front-end and the back-end via copper cables of up to 1km. The DC BCT uses slow sampling ADCs. The data rate is compatible with the transmission bandwidth and every sample is transmitted to the surface. The AWAKE BPM system samples at 40Msps. The data rate would be too high for the transmission of all the samples to the surface. For this reason the front-end implements a threshold based auto triggering algorithm to implement data reduction. The relevant data is buffered and transmitted with a 1Msps protocol taking advantage of the very low repetition rate of the experiment. Both those systems use custom, non-standard, digitizer front-end systems and FMC mezzanines on the VFC-HD.

The Beam Loss Monitoring (BLM) system based on diamond detectors as well as the LHC and SPS fast BCT systems will use a fast sampling commercial ADC FMC mezzanines, the FMC1000 from Innovative, plugged directly on the VFC-HD for their acquisition systems. A similar solution is under study for the new LHC interlocked BPM system, where the sampling speed will be most probably at 4Gsps.

The standard LHC BLM system has its digitalization modules in radiation exposed areas, but at the design time the GEFE was not yet specified. The front-end of this system integrates in the same board the analogue and digital part. The structure of the digital section of such system as well as its form factor is similar to that of the GEFE. It makes use of an anti-fuse FPGA from Actel to implement the digitalization control logic and the interface to the optical link based on the Gigabit Optical Link (GOL), the GBT predecessor with a 640Mbps payload bandwidth. This front end has been tested and qualified including the firmware for a total dose of 70krad. Two of such front-ends with redundant links will connect to a single VFC-HD in a future upgrade of the back end part of the system.

The new Multi Orbit POsition System (MOPOS) of the SPS will be the first instrument to use both the VFC-HD and the GEFE modules. The digitalization is performed on a custom mezzanine equipped with two 4 channel rad-hard 12bit ADCs @10Msps. The sampling clock will be derived from the beam synchronous one recovered by the GBTx. All the samples will be sent interleaved from the GEFE to the VFC-HD where they will be received as if the ADC would be on board thanks to the latency determinism of the link.

REFERENCES

[1] A. Boccardi *et al*, "Ongoing electronic development in the CERN Beam Instrumentation Group: challenges and solutions for the measurement of particle accelerator beam parameters", TWEPP2012, Oxford, United Kingdom, September 2012

[2] The White Rabbit project page: http://www.ohwr.org/projects/white-rabbit

[3] The GBT project web page: https://espace.cern.ch/GBT-Project/default.aspx

[4] M. Barros Marin et al, "The GBT-based Expandable Front-End (GEFE)", TWEPP2015, Lisbon, Portugal, September 2015

[5] The FEASTMP (DCDC converter) project web page: http://project-dcdc.web.cern.ch/project-dcdc/