ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TUP044

THE DIGITAL SIGNAL PROCESSING CHAIN OF THE CERN LIU BWS

D. Belohrad*, J. Emery, J. C. Felipe,
A. Goldblatt, A. Guerrero, M. M. Nieto, F. Roncarolo
European Organisation for the Nuclear Research (CERN), Geneva, Switzerland

Abstract

Between 2019 and 2023, as part of the LHC Injectors Upgrade (LIU), a major renovation of the CERN wire scanners (BWSs) was performed. The main driving force was to prepare the wire scanners for the High-Luminosity LHC (HL-LHC), during which the instantaneous luminosity is expected to double, to around $5 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$. In 2021 seventeen LIU BWSs were installed in the CERN PS complex and the Super Proton Synchrotron (SPS). Additionally, two BWSs were installed in the LHC, at the end of 2022, to be ready for the 2023 LHC run. The contribution aims to describe in detail the technical implementation of the data acquisition chain (DAQ) and digital signal processing of the photomultipliers of the newly installed BWSs. Particular attention is given to the design of the analogue front-end, signal conversion, and data processing chain — providing data for the profile reconstruction. The synchronisation of the incoming digitised signal with the machine timing is also addressed, as it differs significantly between the PS complex and the LHC and SPS. Finally, the limitations of the system are discussed.

INTRODUCTION

The BWS functional and engineering specifications, as well as commissioning and first operational results, were already documented in the last few years [1–3]. This paper is intended to complement the information about the LIU BWS project with the details about the new DAQ design, development, and implementation.

The previous BWS DAQ, completed in the early 2000s, was based on in-house developed digital acquisition board (DAB), equipped with two individual bunch measurement systems (IBMSs) [4, 5]. They were general-purpose platforms targeting various LHC and SPS instrumentation.

The IBMSs included an analogue integrator ASICs, able to deliver bunch-per-bunch information, e.g. of the secondary particle shower used to reconstruct beam profiles with BWSs. These were converted into digital data streams using 40 MSPS 14-bits ADCs, and then processed by the DAB's FPGA.

As discussed more in detail later in the paper, a new DAQ generation was designed and implemented, as part of the LIU BWS project. The DAQ is based on a modern FPGA VME FMC Carrier [6], equipped with a fast FMC ADC, which allows the bunch-per-bunch signals reconstruction via digital integration.

* david.belohrad@cern.ch

DAQ CHALLENGES AND SOLUTIONS

The LIU BWSs installed in the entire CERN accelerator complex use the same DAQ. This makes maintenance and spares management easier, but implies challenges in the DAQ gate-ware design to adapt to the different accelerators beam parameters. In order to set up the systems, diagnose problems and function during normal operation, the DAQ has to produce bunch-per-bunch integrals together with raw data acquired by the ADC.

The differences between the Proton Synchrotron Complex (CPS) and SPS installations are formalized as follows.

In the CPS the beam can be represented as:

- *bunch pattern* in the Proton Synchrotron (PS), and
- *continuous flow of bunches*, as e.g. in the Proton Synchrotron Booster (PSB), (Fig. 1).

In the CPS the bunches are separated into their respective turns by *Turn tag*s. The turn tag is a revolution period signal delivered by the accelerator RF subsystem. The number of bunches in the turn is supplied by SW as a *harmonic number* (H). Knowing the total number of samples in turn (S_{tot}) and H we determine the *integrating window* for each bunch. The bunch-per-bunch integrals $S_{bunch,turn}$ are calculated as the sum of samples over the integrating windows.

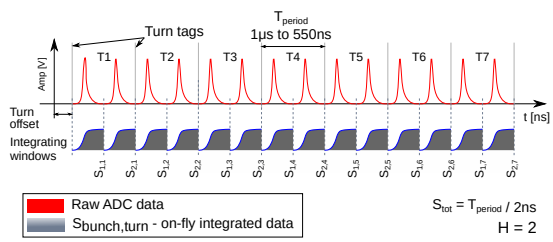

Figure 1: Example of an ideal PSB data acquisition, assuming there are two bunches and the wire passes through a transverse portion of the beam with a constant number of particles during the 7 turns displayed.

In the SPS the integrating windows are not artificially created from the harmonic number. They are identified by the bunch tags delivered together with turn tags via the beam synchronous timing (BST) [7]. The bunch-per-bunch integrals are calculated as the sum of samples between two consecutive bunch tags.

In both cases, a major challenge lies in tracking the tags, since their period changes during the acceleration cycle. In the worst case of the CERN PSB the period changes from 1 µs at injection to ≈ 550 ns at extraction. Another difficulty is to find a correct phase relation between the tags delivered to DAQ and the bunches in the sampled data. The tags

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TUP044

have to fit to 'no-beam' in the data to minimise cross-talk between consecutive bunch slot integrals. Assuming the bunch spacing (used for the bunch tags) is well known, the phasing relies on the *Turn offset* shown in Fig. 1. Both the turn offset and the number of ADC samples per bunch must be set before starting the ADC data processing. The DAQ trigger (occurring once per scan) is not synchronised with the machines' timing (defining the turn clocks) and RF settings (defining the first bunch position in the turn and the bunch spacing). Therefore, the turn offset and the bunch tags setting cannot rely on previous acquisition sets. In addition, as the acquisition starts when the wire does not yet intercept the beam, the first turn signal cannot be used either for setting on the fly the phasing. Therefore, the DAQ features two processing modes that are implemented in the same way for the CPS and SPS systems:

1. **Capture**: The ADC samples are tagged by turn and bunch marks and written into volatile memory, together with S_{tot} per turn.
2. **Integration**: The volatile memory is read back to calculate the bunch integrals.

The raw data from up to 64 scans can be captured and stored in the volatile memory. It is used to estimate the correct turn offset, perform the digital integration and reconstruct the beam profile from any of the ADC channels, as many times as needed.

DAQ IMPLEMENTATION

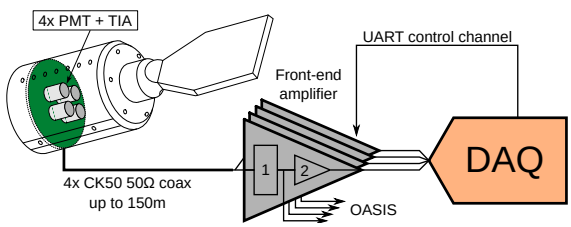

Figure 2: BWS data acquisition chain.

The acquisition chain of a BWS is illustrated in Fig. 2. The secondary particle shower, produced when the BWS wire interacts with the beam, is intercepted by a scintillator coupled to 4 separate photo multipliers (PMTs) (Hamamatsu R9880U). They share the same HV power supply and are read-out in parallel. In order to enhance the detection dynamic range at a constant PMT voltage, optical density filters of different attenuation are installed between the scintillator and the 4 PMTs. For each machine and its beam types, a common PMT voltage is defined such, that at least one of the ADC channels is not saturated and provides enough signal to noise ratio (SNR).

Signal Conditioning

To maximise the SNR of the measurement, the PMTs signals are amplified directly at the source. Current-feedback *trans-impedance amplifier*s (TIAs) with a gain of 500 mV/mA and output impedance of 50 Ω convert the PMT current-source signals to voltage. Four coaxial cables bring the analogue signals to the surface *front-end amplifier* (FEA) that provides a two-stage amplification and filtering to the DAQ input. In the first stage, the signals pass through 260 MHz non-reflective low pass filters, are buffered and sent in parallel to the CERN OASIS monitoring system [8], and to the second FEA gain stage. The gain setting for the second stage, switchable between 0 dB and 10 dB, is common for all four channels. Each second gain stage is also equipped with adjustable bipolar offset used to minimise the DC baseline shifts. Antialiasing filters, with a cut-off frequency of approximately 250 MHz, are used to remove high-frequency content. The offset and gain of the FEA can be set remotely via SW through the *UART control channel*.

The analogue bandwidth of both amplifiers exceeds the system functional specification of 100 MHz. The overall bandwidth limitation comes from the coaxial cables used (type CERN CK50, ≈ 6 dB/100 m at 400 MHz). In some BWS installations, the coaxial cable lengths exceed 150 m. This causes the analogue signal pile-up at the DAQ inputs such, that the individual bunches cannot be clearly separated. Such a *cross-talk* has to be minimised by the digital signal processing.

Analogue to Digital Converter

The DAQ for the CERN LIU BWSs is composed of a new generation VME FMC Carrier (VFC) [6] developed at CERN, hosting an ADC FMC, with specifications shown in Table 1.

Table 1: BWS ADC Specifications. ENOB and SFDR Compliance is Required in the Bandwidth of 70 — 250 MHz

Property	Value
Number of channels	4
Sampling rate	≥ 500 MSPS
Digital interface	JESD204B
Form factor	FMC, VFC compatible
Analogue input coupling	AC, DC
ENOB	≥ 9.5
SFDR	≥ 70 dBc

In 2019, it was not possible to find FMCs of this type off-the-shelf, so CERN conducted a market survey to procure a prototype and then a series of 60 pieces. The tender was won by IAM Electronics GmbH. In October 2020, there were thirty 4-channel 14-bits 500 MSPS FMC ADC modules available at CERN. The design was released under the CERN Open Hardware Licence [9].

The block schematic of the module is shown in Fig. 3. The analogue input consists of four DC-coupled single-ended to differential amplifiers (LMH5401), followed by passive 900 MHz LC filters. Those define the analogue input bandwidth of the sampling system. The analogue signals from

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TUP044

each channel are fed into a single 500 MSPS 4-channel ADC from Texas Instruments. The converted digital data are streamed through high-speed serial lines to the FPGA using the JESD204B protocol [10]. The JESD204B synchronisation signals are generated by the FPGA gate-ware (shown in orange in Fig. 3).

The fast ADC sampling rate and high-speed serial transmission require ultra-clean clock with sub-picosecond jitter. Such clock is delivered by jitter-attenuating PLL from Silicon Labs. An external 48 MHz crystal is used in the loop-back of the internal PLL oscillator to define the base frequency. The PLL outputs provide not only 500 MHz reference for the ADC, but also the 1.016 MHz SYSREF for the JESD204B, and a 500 MHz clock for the FPGA transceiver block (shown in green in Fig. 3). To reduce even further the SFDR, the installed DC/DC converters synchronise their switching to SYSREF.

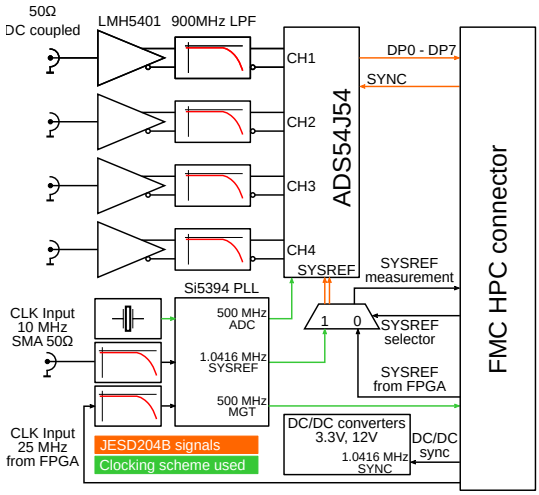

Figure 3: Simplified block schematic of the 4-channel 14-bits 500 MSPS FMC ADC module from IAM Electronics GmbH.

Firmware Implementation

The VFC is a VME64x carrier chosen to house the FMC ADC. The carrier was developed and produced specifically for the CERN Beam Instrumentation group to serve as a common development platform. It is equipped with a single FMC high-pin count slot, two 8 Gbit DDR SDRAMs, and a 115k-cell Intel ArriaV GX FPGA. The board can receive the SPS and LHC BST either through an optical interface or through the P0 VME connector when installed in a LHC type VME crate. The DAQ raw data signal processing is entirely implemented in the VFC FPGA. The processing chain is shown in Fig. 4.

The ADC data stream is received through eight 5 Gbit transceiver lanes (HSSI). As the transceivers are *not* located on the same side of the FPGA, the JESD204B subsystem has to be instantiated manually — by generating two physical interfaces (PHYs), using two different clock sources, joined by a JESD204B MAC. This creates a mesochronous clock network in which the (unknown) phase between the clocks is defined by physical routing and does not represent a problem.

The JESD204B IP core delivers data as a stream of 256-bit frames. The frame clock frequency is 125 MHz, so a data frame is received every 8 ns. Each frame contains 4 samples of 4 channels. The ADC samples, which are 14 bits, are left-aligned in 16-bit bit vectors. The two least significant bits (LSBs) of each sample are set to zero.

In capture mode, the channel samples are tagged to identify where bunches and turns start within the data stream. For each channel, the corresponding sample's LSB within a frame is set when tags are intercepted in the timing signal. For this to work properly, the SPS timing signal has to be *reconstructed* in the 500 MHz ADC clock domain. In the original 160 MHz BST clock domain, the turn tags always come in phase with bunch tags. Thus, the turn tag automatically identifies the start of the bunch as well. Due to a meta-stability caused by the clock domain crossing, the tags can be — in the ADC clock domain — mutually delayed by a single clock cycle, in any order. The turn tag has to be re-aligned in the ADC clock domain to the closest bunch tag received.

Two DDR SDRAM IP cores, clocked at 533 MHz, communicate with SDRAMs. This clock speed was chosen as a compromise between achieving maximum transfer speeds and minimising failures to close the timing of the FPGA. For such configuration, the SDRAM connecting interface is clocked at 125 MHz. Sharing the same clock speeds between the JESD204B and DDR controllers for storing the tagged samples into the SDRAMs would be advantageous. Unfortunately, the VFC SDRAMs do not provide enough transfer bandwidth, as it is significantly lowered by SDRAM refresh cycles. The only possibility to write the tagged data into the SDRAMs is to reduce their amount and provide an elastic buffer to cover the dead time. This means storing only 11 bits out of 14 of each sample, plus one control bit. By capturing 4 channels at the same time, we can encode the bunch and turn marks into control bits of 2 different channels, and use the other 2 remaining channels to serially encode S_{tot}. The 12-bit datagrams are packed into two 128-bit data streams by the *data packing* block shown in Fig. 4. Packed data is sent to elastic buffers to be scheduled for writing into the SDRAMs. The size of the elastic buffers was designed to cover a 20 ms acquisition window, using 40 % of the available internal FPGA memory. Capture can be triggered by both SW and hardware (HW). In the real environment, the trigger comes through the HW from the BWS resolver [11]. The rising and falling edges of the trigger signal determine the start and stop of a single acquisition. Up to 64 consecutive acquisitions can be stored in the SDRAM, and the *index memory* identifies their start and end addresses. During the capture, the amplitudes of all 4 channels in the digitised signals are inspected. The amplitudes are sorted into 8 levels, and the maximum per-channel levels reached are encoded together with the addresses into the 64-bit index memory record. The SW can efficiently decide on the best channel to use for the profile generation by analysing this information.

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TUP044

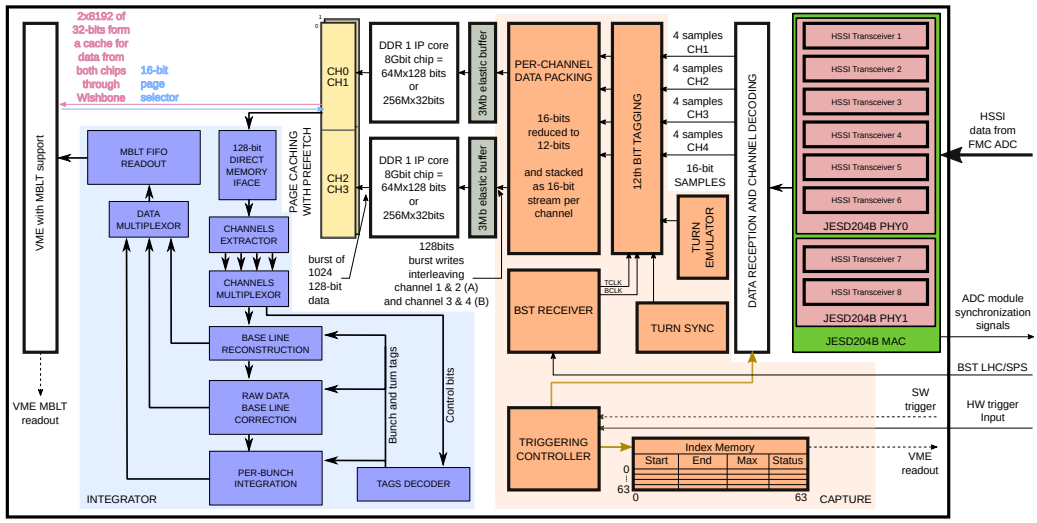

Figure 4: FPGA design block schematic.

In the integrating mode, either raw ADC or per-bunch integrated data can be obtained. The signal processing is applied to the data *previously captured, and stored into the SDRAMs*. Raw data serve mainly to identify the turn offset. The integrated data are used to reconstruct the profile.

Once a chunk of raw data is read, the offset is determined by searching for the delay between the bunch tags and the bunch apexes. This information — together with the SDRAM address start and the best channel to use — is written by the SW into the integrator. By starting the integrator, the *direct memory interface* (DMI) (Fig. 4) pre-fetches and continuously delivers the pages of the captured data from the SDRAMs. The *channels extractor* unpacks the 128-bit data back into 4 separate channels and their respective control bits. The *tags decoder* reconstructs from the samples' control bits the tags and S_{tot}. In case of CPS BWSs the bunch tags are generated *on-the-fly*. The tags are used in successive blocks to manipulate the raw data. Only one raw data channel at a time can be selected for further processing. Figure 5 shows an example of the raw and integrated data, as well as the tags, of the uncorrected acquisition.

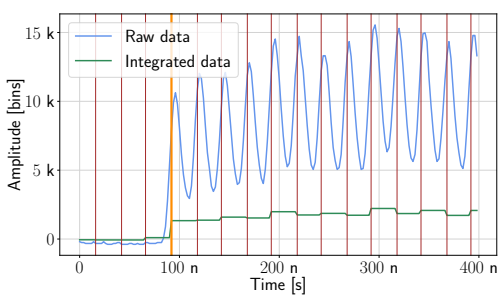

Figure 5: Raw and uncorrected PS acquisition on `PR.BWS.65H`, LHC cycle with 48 bunches (H=84), no baseline correction.

The bunch tags together with the selected raw data are used to regenerate the baseline of the signal. The baseline data stream is constructed by sampling raw data at two consecutive tags, and generating the linear approximation between them. Such a stream is subtracted from the original data to suppress the baseline and therefore minimise the signal pile-up. The *Per-Bunch Integrator* sums all the corrected data between two consecutive bunch tags. A VME A32D64 multiplexed block transfer (MBLT) is used to deliver the data to the CPU, achieving the transfer speed of ≈ 30 MBytes/s.

A card register allows for the disabling of baseline restoration and integration. This can be helpful, for example, for reading raw ADC data through the MBLT mechanism. Additionally, each processing block can be individually connected to the VME interface to inspect intermediate signal processing products.

CONCLUSION

The article presents the DAQ for the BWSs installed in the entire CERN accelerator complex. Since January 2021, more than **150 thousands scans** were performed and the DAQ is fully commissioned for the operation in the CPS and SPS. The commissioning of the DAQ for the LHC BWSs is expected in 2024.

Using VFC in combination with FMC ADC from IAM Electronics represents a trade-off between fully exploiting the technical capabilities of the ADC module, and the possibility to use an in-house common platform for development. While the presented solution fully complies with the required BWS specifications, it can be limiting for other CERN beam instrumentation systems, where 4-channel raw data need to be stored into the VFC DDR SDRAMs.

REFERENCES

[1] J. Emery *et al.*, "Commissioning of the LHC injectors BWS upgrade," in *Proc. IBIC'22*, Pohang, Korea, Sep. 2021, pp. 239–243. doi:10.18429/JACoW-IBIC2021-TUPP18

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TUP044

[2] A. Guerrero, D. Belohrad, J. Emery, S. Jackson and F. Roncarolo, "Modular software architecture for the new CERN injector wire-scanners", in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 487-491. doi:10.18429/JACoW-ICALEPCS2021-TUPV037

[3] S. Di Carlo *et al.*, "Commissioning of the LHC injectors upgrade fast wire scanners and first experimental results", *Nucl. Instrum. Methods Phys. Res.*, vol. 1053, p. 168328, 2023. doi:10.1016/j.nima.2023.168328

[4] E. Calvo *et al.*, "The LHC orbit and trajectory system", in *Proc. 6th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators (DIPAC'03)*, Mainz, Germany, May 2003, pp. 187–191.

[5] A. Guerrero, H. Jakob, R. Jones, and J. J. Savioz, "The SPS individual bunch measurement system", in *Proc. 5th European Workshop on Diagnostics and Beam Instrumenation (DIPAC'01)* Grenoble, France, May 2001.

[6] A. Boccardi, "Ongoing electronic development in the CERN Beam Instrumentation Group: challenges and solutions for the measurement of particle accelerator beam parameters", *J. Instrum.* vol. 8, Mar. 2013. doi:10.1088/1748-0221/8/03/C03006

[7] J.J. Savioz, "BOBR: The beam synchronous timing receiver interface for the beam observations", Technical Specification, CERN, Geneva, Switzerland, Rep. LHC-BI-ES-XXXX.XX rev 1.0 Draft, http://www.cern.ch/TTC/BOBRspec.pdf

[8] S. Deghaye, D. Jacquet, I. Kozsar, and J. Serrano, "Oasis: A new system to acquire and display the analogue signals for LHC", in *Proc. ICALEPCS'03*, Gyeongju, Korea, Oct. 2003.

[9] Open Hardware Repository, "FMC ADC 500M 14b 4CHA", https://ohwr.org/project/fmc-adc-500m14b4cha

[10] Texas Instruments, "JESD204B Overview", https://www.ti.com/lit/pdf/slap161

[11] J. Emery and F. Roncarolo, "CERN LIU beam wire scanner behavioural interface between hardware-firmware-software", Engineering Specification, CERN, Geneva, Switzerland, Rep. EDMS:2665438, https://edms.cern.ch/document/2665438