ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPA093

SPS BUNCH-BY-BUNCH PHASE MEASUREMENT IN THE CERN SPS LOW LEVEL RF

R. Borner*, P. Baudrenghien, A. Spierer
CERN, Geneva, Switzerland

Abstract

As part of the High Luminosity LHC (HL-LHC) project, the Super Proton Synchrotron (SPS) Low Level RF (LLRF) system has undergone a complete re-design. During multi batch injection into the machine we must be able to distinguish between bunches already circulating in the machine and newly injected bunches. This requires a system that can measure the phase of each individual bunch circulating in the SPS, where the bunch spacing is 5 ns. This is to be used for both diagnostics and as an input to the Beam-Based phase loop where it will also be required to properly operate two phase loops during slip stacking. To facilitate this, an FMC ADC mezzanine card was chosen with a sampling rate of up to 6.4 Gsps. This was paired with the AFCZ µTCA carrier board which uses a System on Chip (SOC) FPGA on which the signal processing of the data is carried out. The design of the FPGA firmware presented some interesting challenges as the signal processing algorithms must process many samples in parallel due to the high throughput of data. In this case the ADC sampling clock is twenty times faster than the FPGA processing clock. The paper presents the motivations for the upgrade, the overall architecture, the main algorithms, and details about the hardware and firmware implementation.

MOTIVATION

Bunch per bunch measurements have been very useful for diagnostics in the LHC. Longitudinal measurements (bunch phase) are used to monitor machine scrubbing [1] and to study coupled-bunch instabilities. A similar tool would also help to optimize the SPS. In addition, these could be used in the operational Beam Based Loop. Since its restart in Spring 2021 the SPS phase loop uses signals from a resonant PU, resulting in an averaging over 500 ns [2]. This has proven sufficient for the proton beams, but it does not give clearly separated measurements during multi-batch injection from the PS, with 225 ns batch spacing. In addition, the new production of LHC ions beams calls for the momentum slip stacking of two families of bunches [3]. During this manipulation, each family has a phase loop considering only the bunches of the family, and disregarding locations where the two families overlap. This calls for a bunch per bunch phase measurement.

ALGORITHMS

Short Time Fourier Transform

We start with the data stream x_n at 5 Gsps sampled from a Wideband longitudinal pickup (analog bandwidth in excess of 2 GHz). In the SPS we can consider an RF frequency that is between 199.9 and 200.4 MHz. With a sampling frequency of 5 Gsps, we therefore have maximum of 26 samples in any given bucket. The RF component of the beam current from that bucket (amplitude and phase) is calculated from the Short Time Fourier Transform (STFT)

$$X = \sum_{n=bucketStart}^{bucketStart+N-1} x_n e^{-j\varphi_n^{RF}}. \tag{1}$$

The SPS LLRF broadcasts the instantaneous frequency as a numerical Frequency Tuning Word (FTW) on the White Rabbit (WR) link [2]. The FTW is inputted to a local Numerically Controlled Oscillator (NCO) that reconstructs the RF phase. Successive samples are padded with zeros such that we have N=32 samples to send to the Fourier transform processing block. In Eq.(1) the time index is 5 GHz but the phase of the RF is available from the NCO at 125 MHz [2]. We therefore need to interpolate 40 RF phase samples between the available NCO samples. Let φ_0^{RF} be the RF phase at the beginning of the bucket and $\Omega_0^{RF,NCO}$ be the RF frequency normalized to 5 GHz, we have

$$X = e^{-j\varphi_0^{RF}} \sum_{n=0}^{N-1} x_n W^n, \tag{2}$$

where

$$W = e^{-j2\pi\Omega_0^{RF,NCO}}. \tag{3}$$

The processing of this equation can be efficiently implemented using a variant of the butterfly computation (FFT) [4]. We can use a Flow-Graph to represent the processing. All intermediate data steps are complex values. Fig. 1 shows the tree structure that implements the algorithm Eq.(2). The nodes represent registers in the FPGA implementation. Processing goes from left to right. If a branch has a label, it implies multiplication by that value. Two branches merging into a node implies an addition. On the left we have the data x_n: x_0 is the first data sample after a detected bucket start. The data record extends to the end of the bucket and is then padded with zeros to 32 values. These data samples are arranged in bit-reversed order. On the top left we also input the phase and frequency given by the RF NCO at the time of the first data sample x_0. The processing contains six successive stages. The RF NCO phase and frequency must

[*] robert.borner@cern.ch

be pipelined so that it remains synchronized with the data x_n. With the parallel processing algorithm shown on Fig. 1 we can process the 26 data acquired at 5 Gsps at the 250 MHz FPGA clock rate, allowing for one measurement per bucket.

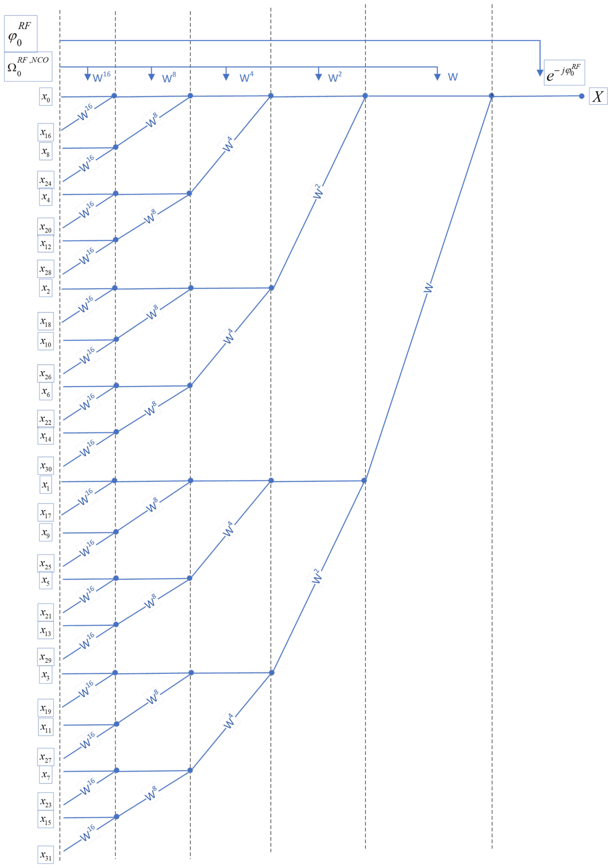

Figure 1: Butterfly algorithm.

Data Stream Slicing into Successive RF Buckets

The 5 Gsps data stream from the FMC217 mezzanine card is received continuously by the FPGA in batches of 20 parallel ADC samples at a rate of 250 MHz. It is therefore necessary to slice the data into successive buckets and align the data such that the first sample of each bucket can be aligned with the first input sample of the Fourier transform. To find the bucket start we must find the point at which the RF phase crosses zero with a granularity of 200 ps (5 Gsps) but the RF phase is only updated from the numerically controlled oscillator (NCO) at a rate of 125 MHz. We therefore need to interpolate 40 RF phase samples in between phase updates from the NCO. For interpolation of the RF phase we use the $\Omega^{RF,NCO}$ encoded in the FTW.

As the phase is represented as unsigned with 2π as full scale, a bucket start or end corresponds to a transition from a large value to a small value and can easily be detected as a change in the most significant bit (MSB) from 1 to 0. The ADC provides 20 data samples at a rate of 250 MHz, while the NCO phase and FTW are available at only 125 MHz. On each 250 MHz cycle we compute in parallel the 21 RF phases, using the last available NCO phase and FTW. Comparing the successive RF phases we detect if the MSB has changed. If so, there has been a bucket transition between the corresponding two 5 GHz data samples.

HARDWARE

The Beam Phase module has been designed on the uTCA platform. The beam phase module is based on an AFCZ mezzanine carrier with a Xilinx Zynq Ultrasacle+ FPGA and the Vadatech FMC217 mezzanine card (see Fig. 2). The variant of the Vadatech FMC217 ADC FMC that was chosen can be configured to run at sample rates of up to 6.2 Gsps. For this application the FMC217 has been configured to run at a sample rate of 5 Gsps. This sample rate was chosen because the frequency of the main processing clock in the FPGA, which is synchronised to the white rabbit core, has a frequency of 250 MHz. Using a sample rate of 5 Gsps therefore gives us an even integer ratio of 20 ADC sample periods per FPGA processing clock period.

Figure 2: AFCZ with FMC217.

FIRMWARE

Processing Architecture

The bunch by bunch phase processing architecture is shown in Fig. 3. The ADC FMC samples the signal from the beam pickup at a rate of 5 Gsps. The ADC clock is locked to a 250 MHz white rabbit clock input. The ADC interface parallelizes this data such that it outputs an array of 20 ADC samples at a rate of 250 MHz. The NCO outputs the FTW and H1 (revolution frequency) phase at a rate of 125 MHz. The FTW and H1 phase are read by the phase zero crossing detection block which calculates the phase advance in 5 Gsps increments for the next 20 samples. The position at which the H4620 phase (4620 is the SPS harmonic number at 200 MHz) crosses zero is sent to the FFT sample selector block which is responsible for slicing and aligning the continuous data stream such that the ADC sample at which the phase crossed zero is inputted into the first position of the FFT. Successive samples are then inputted in the correct order. The Butterfly algorithm FFT is implemented as a tree of complex add multiply elements as described in Fig. 1. The butterfly algorithm has a delay of 54 clock cycles, this in itself is not an issue as the data input is pipelined, but all

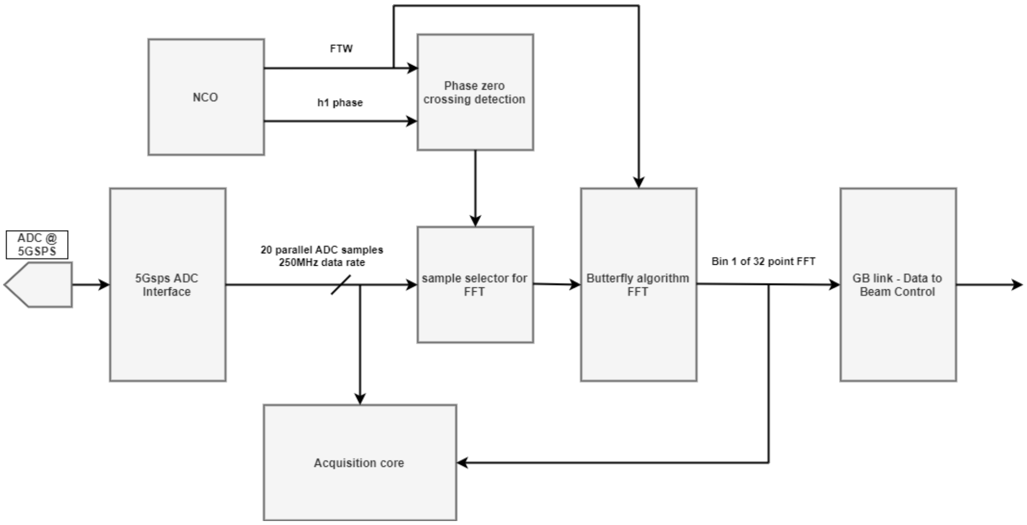

Figure 3: Processing architecture.

other status and data signals must be delayed by the same amount to ensure that all paths through the design have the same latency. All valid results from the FFT are identified with a valid flag, these valid results are used by both the acquisition core where they can be used for diagnostics, and sent via the GBLink to the Beam Control module to be used in the the beam based loops.

Data Acquisition

The ADC used for the bunch by bunch phase measurement is configured with a sampling rate of 5 Gsps. This is a higher clock speed than can be supported by the FPGA. For this reason, the ADC interface takes the incoming ADC data stream and parallelizes the data so that the data can be received at a rate supported by the FPGA. In the beam phase module the data is received in an array of 20 ADC samples at a data rate of 250 MHz. For setting up and diagnostics, being able to properly acquire and display this data is necessary.

An acquisition core has been developed at CERN [5] for other modules on the µTCA platform. This acquisition core was implemented and configured with 20 channels; one for each position in the array. With this configuration each channel will receive 1/20th of the data. To display this data properly a modified version of the acquisition core software was developed to read each of the 20 buffers, interleave the data in time order and display it as a single waveform.

The standard data width of the acquisition core IP is 32 bits, and the ADC is 12 bit, the memory space can further be optimised by combining two ADC samples in each memory location. This way the required number of channels has been reduced to 10.

RESULTS

The firmware full architecture and algorithms have been implemented on the FPGA. The algorithms have been tested in simulation and shown to perform as expected. Setup scripts for the complex ADC interface have been migrated to run on the embedded ARM core inside the Zynq Ultrascale FPGA. The ADC interface and acquisition has been tested on the hardware in the lab test setup and along with the acquisition software that has been developed, we are able to acquire and display properly the 5 Gsps ADC data. The next steps are to test the full processing chain on the hardware on the lab test setup. Once this has been completed testing and commissioning will commence in the SPS.

REFERENCES

[1] J. F. Esteban Müller, P. Baudrenghien, T. Mastoridis, E. Shaposhnikova, and D. Valuch, “High-accuracy diagnostic tool for electron cloud observation in the LHC based on synchronous phase measurements”, *Phys. Rev. ST Accel. Beams*, vol. 18, 2015, p. 112801.
`doi:10.1103/PhysRevSTAB.18.112801`

[2] A. Spierer *et al.*, “The CERN SPS Low Level RF: The Beam-Control”, in *Proc. IPAC’22*, Bangkok, Thailand, Jun. 2022, pp. 895–898. `doi:10.18429/JACoW-IPAC2022-TUPOST021`

[3] P. Baudrenghien, J. Egli, G. Hagmann, A. Spierer, and T. W, “The CERN SPS Low Level RF: Lead Ions Acceleration”, in *Proc. IPAC’22*, Bangkok, Thailand, Jun. 2022, pp. 899–902. `doi:10.18429/JACoW-IPAC2022-TUPOST022`

[4] J.W. Cooley and J.W. Tukey, “An Algorithm for the Machine Calculation of Complex Fourier Series”, *Math. Computation*, vol. 19, 1965, pp. 297-301.

[5] J. Egli, A. Spierer, G. Hagmann, M. Suminski, and P. Baudrenghien, “The CERN SPS Low Level RF: embedded acquisition system for the Cavity-Controller and Beam-Control commissioning and diagnostics”, presented at the IPAC’23, Venice, Italy, May 2023, paper THPA095, this conference.