

# SIGNAL PROCESSING ARCHITECTURE FOR THE HL-LHC INTERACTION REGION BPMs

D. R. Bett\*, University of Oxford, Oxford, UK

A. Boccardi, I. Degl'Innocenti, M. Krupa, CERN, Geneva, Switzerland

## Abstract

In the HL-LHC era, the Interaction Regions around the ATLAS and CMS experiments will be equipped with 24 new Beam Position Monitors (BPM) measuring both counter-propagating beams in a common vacuum chamber. Numerical simulations proved that, despite using new high-directivity stripline BPMs, the required measurement accuracy cannot be guaranteed without bunch-by-bunch disentanglement of the signals induced by both beams. This contribution presents the proposed signal processing architecture, based on direct digitisation of RF waveforms, which optimises the necessary computing resources without a significant reduction of the measurement accuracy. To minimise the number of operations performed on a bunch-by-bunch basis in the FPGA, some of the processing takes place in the CPU using averaged data.

## INTRODUCTION

The Large Hadron Collider (LHC) will undergo major upgrades in the context of the High Luminosity LHC (HL-LHC) project with the goal to deliver  $3000 \text{ fb}^{-1}$  of integrated luminosity over twelve years of operation from 2027 [1]. New Inner Triplets (IT) consisting of several high-gradient focusing magnets around ATLAS and CMS experiments will squeeze the proton beams to a  $7.1 \mu\text{m}$  RMS beam size at the collision point [2]. In order to reliably collide such exceptionally small beams, each HL-LHC IT will feature six Beam Position Monitors (BPM) of two different types [3]. Since these BPMs will be installed in regions where both proton beams circulate in a common vacuum chamber, they must be able to clearly distinguish between the positions of the two counter-propagating particle beams.

The longitudinal positions of the BPMs were optimised to guarantee that the temporal separation between the two beams at each BPM location will always be greater than 3.9 ns, which is approximately 3 times longer than the bunch length. Nevertheless, using directional-coupler BPMs (also known as stripline BPMs) is unavoidable to reduce the inter-beam cross-talk. In such BPMs the passing beam couples to four long stripline electrodes parallel to the beam axis. Each electrode is connectorised on both ends but the beam couples predominantly to the upstream port with only a relatively small signal generated at the downstream port. This feature, referred to as directivity, allows both beams to be measured by a single array of electrodes. Figure 1 shows a 3D model of one of the HL-LHC stripline BPMs. Most of the HL-LHC IT BPMs incorporate four tungsten absorbers protecting the superconducting magnets from the high-energy collision




Figure 1: Tungsten-shielded cryogenic directional coupler BPM design for HL-LHC.

debris [4]. As the absorbers must be placed in the horizontal and vertical plane, the BPM electrodes are installed at  $\pm 45^\circ$  and  $\pm 135^\circ$  significantly increasing the measurement non-linearity for large beam offsets.

To cope with the very demanding requirements of precise beam position measurements near the experiments, a new state-of-the-art acquisition system for the HL-LHC IT BPMs is under development. It will be based on nearly-direct digitization by an RF System-on-Chip (RFSoc) [5]. This unique family of integrated circuits combines a set of Analogue-to-Digital Converters (ADC), Digital-to-Analogue Converters (DAC), Programmable Logic (PL) and several embedded CPUs, referred to as the Processing System (PS), on a single die. Each of the 8 ports of each BPM will be connected to a dedicated RFSoc 14 bit ADC channel sampling at  $5 \text{ GSa s}^{-1}$ . The acquisition electronics and signal processing software will use this raw data to compute the beam position applying a correction algorithm to minimize the parasitic contribution of the other beam as well as taking into account the BPM rotation, non-linearity and scaling factors.

## ACQUISITION ELECTRONICS DESIGN CRITERIA

The final specification for the HL-LHC IT BPMs is not yet available but some preliminary design criteria have been set to guide the design of the future acquisition electronics.

The HL-LHC beam will consist of up to 2808 bunches spaced by multiples of 25 ns with intensities spanning close to two orders of magnitude from  $5 \times 10^9$  up to  $2.2 \times 10^{11}$  charges. However, for most common operational scenarios it is assumed that the intensity of bunches within the same beam might vary by a factor of four, while the ratio of bunch intensity between the two beams can reach a factor of ten. HL-LHC bunches are not expected to be longer than 1.2 ns ( $4\sigma$ ) but for some special operational modes the BPM system should be able to measure bunches as short as 0.5 ns.

\* douglas.bett@physics.ox.ac.uk

The temporal separation between bunches of the counter-rotating beams is fixed at each BPM location and ranges from  $\pm 3.9$  to  $\pm 10.5$  ns. The beam position measurement range must cover half of the BPM aperture of approximately 120 mm.

The HL-LHC Inner Triplet BPM system will produce data in two main acquisition modes:

- **Trajectory mode:** on-demand bunch-by-bunch, turn-by-turn measurements over a finite number of turns;
- **Orbit mode:** continuous multi-bunch, multi-shot measurements with data-rate reduced through averaging.

Due to the very high sampling rate, synchronizing the ADC clock with the accelerator timing would be virtually impossible. The HL-LHC IT acquisition will therefore use a free-running clock and will detect the bunch windows automatically based on the sampled data.

Until a full specification is available, the following performance goals are considered at this stage of the project:

- **RMS resolution** (in trajectory mode): 15  $\mu\text{m}$ ;
- **Measurement reproducibility** (over 10 hours in orbit mode):  $\pm 5$   $\mu\text{m}$ ;
- **Maximum two-beam disentanglement error** (in orbit mode):  $\pm 20$   $\mu\text{m}$ .

## OVERVIEW OF SIGNAL PROCESSING

A CST Microwave Studio model of the BPM was used to calculate the signals generated on each BPM port by the passage of a single bunch of charged particles. For a Gaussian bunch (width  $\sigma = 0.3$  ns, intensity  $q = 2.3 \times 10^{11}$  particles) travelling along the BPM axis, the signals induced at opposite ends of a single stripline are shown in Fig. 2.

For each stripline, each beam produces an upstream signal at the entrance end of the BPM and a much smaller downstream signal at the other end; by analogy to a directional coupler, these signals will henceforth be referred to as the “coupled” and “isolated” signals respectively. For a given port, the beam that gives rise to the coupled signal is denoted the “main beam”, while the beam travelling in the opposite direction generating the isolated signal is referred to as the “counter beam”. The total signal observed on any given port is thus the sum of the coupled signal due to the main beam and the isolated signal due to the counter beam. The amplitude of each of the two sub-signals varies with the intensity and position of the corresponding beam, and the two sub-signals are separated in time by the bunch crossing timing for that particular BPM. The graphs presented in this paper are the results of a simulation framework developed in GNU Octave [6].

For the case of two beams of equal intensity, the isolated sub-signal accounts for a very small contribution to the total power observed on each port. However, in the worst case scenario, where there is a large imbalance in the intensity of

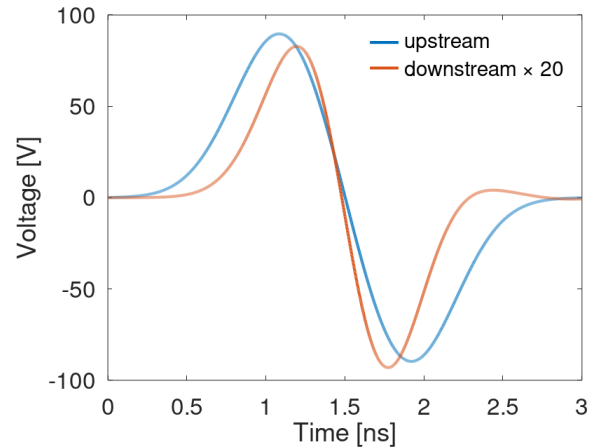


Figure 2: Simulated stripline signals from the CST model of the HL-LHC IT BPMs. The beam enters the upstream end of the BPM. The downstream signal has been amplified by a factor of 20.




Figure 3: Simulated signal on a BPM port for the case of (a) two nominal-intensity ( $q = 2.3 \times 10^{11}$  particles) beams; (b) a main beam with one-tenth of nominal intensity and a counter beam of nominal intensity. The position of the two beams corresponds to the presqueeze orbit of the CP BPM at IR 1 where the separation of the two beams in the vertical direction is more than 25 mm. Note that the main beam arrives first in this case.

the two bunches and a large difference in the orbits, the isolated sub-signal is comparable in magnitude to the coupled sub-signal. This is illustrated in Fig. 3.

### Beam Disentanglement Algorithm

Since the presence of a counter beam may alter the measured position of the main beam, some method of “disentangling” the coupled and isolated sub-signals is required. The method proposed here is named the “power compensation” method as it attempts to subtract the power of the counter beam from the total power measured on each port.

Let's assume the BPM ports being numbered from 1 to 8, with the odd ports being upstream for beam 1 and sharing the stripline with the successively-numbered even port. Let  $V_1[n]$  represent the waveform measured by the digitizer sampling port 1 of the BPM over a single 25 ns period. This waveform can be expressed as:

$$V_1[n] = \kappa_1 V_c[n] + \kappa_2 V_i[n] \quad (1)$$

where:

- $V_c[n]$  is the sequence of samples representing the coupled signal induced by an on-axis reference main beam with some nominal intensity and bunch length;
- $V_i[n]$  is the sequence of samples representing the isolated signal induced by a reference counter beam with the same parameters as the main beam;
- The  $\kappa$  scale factors set the amplitude of each sub-signal according to the actual intensity of each beam and their position in the BPM plane.  $\kappa_1 = \rho_1 \cdot q_1$ , where  $q_1$  is the intensity of the actual beam 1 (expressed here as a multiple of the intensity of the reference main beam) and  $\rho_1$  is a scaling factor accounting for the distance of beam 1 from the electrode with ports 1 and 2.

The signal's "power" can be calculated by taking the sum of squared samples for both sides of the equation:

$$\begin{aligned} \sum_{n=1}^N V_1[n]^2 &= \kappa_1^2 \sum_{n=1}^N V_c[n]^2 + 2\kappa_1\kappa_2 \sum_{n=1}^N V_c[n] \cdot V_i[n] \\ &\quad + \kappa_2^2 \sum_{n=1}^N V_i[n]^2 \end{aligned}$$

where  $N$  is the number of samples in a 25 ns window and  $V_1[n]$  denotes the  $n$ -th sample of the  $V_1$  waveform. By introducing some new notation, the above expression can be written as:

$$\psi_1 = \kappa_1^2 \psi_c + 2\kappa_1\kappa_2 \chi + \kappa_2^2 \psi_i \quad (2)$$

Note that the  $\psi_{c,i}$  and  $\chi$  parameters represent scalars calculated from the observed or reference waveform for the specific port, while the  $\kappa$  parameters are unknowns that we wish to calculate. Equation 2 can be trivially rearranged into the form of an equation quadratic in  $\kappa_1$ :

$$\kappa_1^2 + \left(2\frac{\chi}{\psi_c}\kappa_2\right)\kappa_1 + \left(\frac{\psi_i}{\psi_c}\kappa_2^2 - \frac{\psi_1}{\psi_c}\right) = 0 \quad (3)$$

By making the approximation  $\kappa_2 \approx \sqrt{\frac{\psi_2}{\psi_c}}$ , i.e. considering that the influence of beam 1 on port 2 is negligible for computing  $\kappa_1$ , the coefficients of this quadratic equation can be expressed solely in terms of a set of scalars that can be calculated in advance from the reference waveforms ( $\psi_c$ ,  $\psi_i$ ,  $\chi$ ) and a pair of scalars that must be calculated from the measured waveforms at both ends of a single stripline ( $\psi_1$ ,

$\psi_2$ ). The power compensation procedure thus amounts to calculating the coefficients of a single quadratic equation for each port, and then solving them.

### Full Signal Processing Chain

The power compensation represents a single stage of the five-stage signal processing algorithm for converting a set of eight digitized waveforms to a horizontal and vertical position reading for each of the two beams.

1. Calculation of the sum of squared samples for each waveform, i.e.  $\psi_{1..8}$ .
2. Scaling of  $\psi_{1..8}$ . Variable attenuators will be used to match the signal level on each port to the input range of the ADC; this attenuation must be removed before any further processing takes place. The stability of the signals over a single fill (~10 hours) is also an issue of concern and a scheme where the onboard DACs are used to periodically monitor the symmetry of a given pair of channels is being developed. The measured result of such a self-calibration procedure would also be applied at this stage.
3. Power compensation. The values of  $\psi_{n,n+1}$  calculated from the waveforms from the two ports of a single stripline are combined with the  $\psi_{c,i}$  and  $\chi$  values, pre-calculated from the reference waveforms for the specific port, to give the coefficients of a pair of quadratic equations that can be solved to give the amplitude of the main beam signal for each waveform.
4. Position calculation. The usual difference-over-sum method can be used with the main beam amplitudes for each pair of ports (top/bottom and left/right for both ends of the BPM) to give the horizontal and vertical positions of each beam in the BPM frame.
5. A polynomial correction will be applied to the measured horizontal and vertical position of each beam in order to account for the non-linearity and the 45° rotation of the BPM if required.

### Expected Performance

Simulation results suggest that the performance goals will be challenging to achieve given the range of conditions the system must be able to accommodate. While the achieved RMS resolution in trajectory mode consistently works out in the 10-15  $\mu\text{m}$  range, it is reasonable to expect the real system to under-perform relative to the model. To achieve the measurement reproducibility goal in orbit mode, it is calculated that the uncorrected imbalance in signal gain between a pair of ports must not exceed 0.01%. The exact details of the self-calibration scheme are still being worked out, so it is difficult to assess at this stage how feasible such an imbalance is. However, the results suggest that the maximum two-beam disentanglement error of 20  $\mu\text{m}$  can be comfortably achieved using the power compensation procedure.

## SYSTEM ARCHITECTURE

Due to the RFSoc's particular nature, it is possible to distribute the signal processing amongst 3 different systems:

1. The Programmable Logic (PL) of the RFSoc, i.e. the configurable logic blocks, the flip-flops and block RAMs typical of an FPGA;
2. The Processor System (PS) of the RFSoc, i.e. a group of embedded CPUs;
3. The Software (SW) running on a remote back-end computer in communication with the RFSoc.

Even though it would be possible to implement the whole signal processing within the programmable logic of the RFSoc, this would not be the most efficient use of the system.

The acquisition of all signals of one HL-LHC IT BPM requires 8 high-sampling-rate ADCs, one per signal. The PL will be used for fast signal processing. The internal Block RAM will be used for storing high-rate data (raw data for calibration and debugging) and for implementing the filters required for the orbit mode. The external DDR memory will store the high-volume data of the trajectory mode. The PS will control the acquisition, will implement the slow calibration and will be connected to the remote computer e.g. by Ethernet. Orbit mode data will be streamed out continuously at a slow rate (in the order of tens of hertz), while the trajectory mode data will be sent out on demand.

When all signal processing is performed within the RFSoc, the SW receives already calculated beam position data, in arbitrary units, to which it must apply only the calibration coefficients and the non-linearity polynomial correction. Alternatively, the back-end can receive the signal powers computed for each of the 8 BPM ports, after the two-beam compensation and averaging when required, and delegate the beam position calculation to the SW. The latter solution allows to save PL resources, at the cost of doubling the bandwidth and the required memory space. It also provides access to the independent port signals, granting the opportunity for more off-line analyses.

An estimation of the resources needed for the signal processing implementation in both cases is reported in Table 1. Each value is put in relation to the resources available on a Generation-3 RFSoc equipped with 4 GB of DDR4 memory, which are indicated in brackets. An estimation of the maximum rate with which trajectory mode data can be read-out is also given. The PL resources - Digital Signal Processing (DSP) slices, Look Up Tables (LUT) and Flip-Flops (FF) - are accounted only for the implementation of the algorithm and not of the full acquisition system. The available DDR bandwidth (BW) is estimated at about 70% of the declared peak performance in the datasheet of the RFSoc [5]. The required DDR BW is computed for the continuous write of the trajectory data; the read-out is on demand and sporadic. The read-out BW assumes an Ethernet protocol over a 1000 Mbps link running at 50% efficiency. The calculated bandwidth takes into account the back-end continuously reading orbit data from the RFSoc.

Table 1: Summary and comparison of the estimated resources needed when computing the beam position solely in the RFSoc and when a part of the computation is delegated to software (SW) running on a remote computer.

Available resources	Position calculation in RFSoc	Position calculation in SW
ADCs (8)	100%	100%
Block RAM (38 Mb)	3%	4%
PL DSP (4272)	6%	5%
PL LUT (4255280)	3%	2%
PL FF (850560)	3%	2%
DDR (32 Gb)	3%	6%
DDR BW (95 Gbps)	10%	18%
Read-out BW (500 Mbps)	3%	6%
Max trajectory rate	$1.25 \text{ s}^{-1}$	$0.55 \text{ s}^{-1}$

## SUMMARY

Although the signal processing chain for the HL-LHC IT BPMs is still under development, the analysis performed so far indicates that an acquisition system using an RFSoc sampling the 8 BPM outputs could meet the basic performance goals without exceeding the available resources. The decision how to distribute the required processing steps between the RFSoc and the SW running on a back-end computer will be taken based on the experience gained with a proof-of-principle system foreseen for 2022.

## ACKNOWLEDGEMENTS

The authors would like to thank Manfred Wendt from CERN for his continuous input to the project.

## REFERENCES

- [1] L. Rossi, O. Brüning, "Progress with the High Luminosity LHC Project at CERN", in *Proc. IPAC'19*, Melbourne, Australia, May 2019, pp. 17–22. doi:10.18429/JACoW-IPAC2019-MOYPLM3
- [2] F. Bordry *et al.*, "Machine Parameters and Projected Luminosity Performance of Proposed Future Colliders at CERN", [arXiv:1810.13022 [physics.acc-ph]].
- [3] M. Krupa, "Beam Instrumentation and Diagnostics for High Luminosity LHC", in *Proc. IBIC'19*, Malmö, Sweden, Sep. 2019, pp. 1–8. doi:10.18429/JACoW-IBIC2019-MOA002
- [4] L. S. Esposito *et al.*, "FLUKA Energy Deposition Studies for the HL-LHC", in *Proc. IPAC'13*, Shanghai, China, May 2013, paper TUPFI021, pp. 1379–1381.
- [5] Xilinx, DS926 - Zynq UltraScale+ RFSoc Data Sheet: DC and AC Switching Characteristics (v1.8), [https://www.xilinx.com/content/dam/xilinx/support/documentation/data\\_sheets/ds926-zynq-ultrascale-plus-rfsoc.pdf](https://www.xilinx.com/content/dam/xilinx/support/documentation/data_sheets/ds926-zynq-ultrascale-plus-rfsoc.pdf)
- [6] D. R. Bett *et al.*, "Simulation of the Signal Processing for the New Interaction Region BPMs of the High Luminosity LHC", in *Proc. IBIC'20*, Santos, Brazil, Sep. 2020, pp. 120–123. doi:10.18429/JACoW-IBIC2020-WEPP12