A proposal for a trajectory measurement system for the PS Booster

Jeroen Belleman

November 12, 2010

Abstract

This is a proposal to equip the CERN PS Booster with a trajectory measurement system along the same lines as was previously done for the PS. That is, high-speed ADCs convert all BPM signals directly into the digital domain at a high rate, and individual bunch positions as well as averaged orbits are calculated on the fly and stored into a large circular buffer memory. Multiple users may then read the data they are interested in. The system will make use of modern fast ADCs, large FPGAs and SDRAM.

1 Introduction

The Booster is a 1.4 GeV synchrotron consisting of four superimposed rings of 50 m diameter. The beam injected from LINAC 2 is continuous and it is bunched by the accelerating RF system after injection. The harmonic number h can be either 1 or 2. The machine has a normalizer-based orbit measurement system with 16 BPMs per ring, able to measure the orbit of both planes of one of the rings at a time, and which needs about 1 ms from injection to settle.

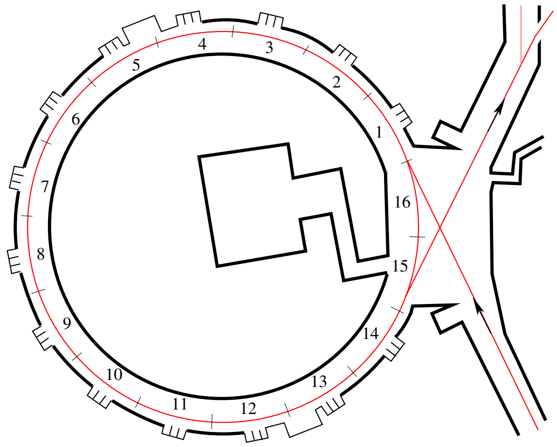

Figure 1: The Booster with injection and ejection lines

A normalizer is a device that continuously regulates the gain of the Σ channel such that its final output level is constant. It then applies the same gain to the X and Y channels, which are thenceforth proportional to beam position only, and no longer to the product of position and intensity. This is a slow process, making orbit acquisitions at less than 1 millisecond intervals pointless. The low bandwidth and inherent automatic gain control lends to normalizers the ability to function over a large range of beam intensities. They are also simple to use, in that they need no tuning and no timing or RF signals. They cannot provide trajectories however, but only averaged orbits.

With the new LINAC 4, the beams injected into the PSB will already be bunched, presenting the opportunity of measuring trajectories of individual bunches right from injection [1].

2 PSB parameters

The Booster revolution frequency f_{rev} for protons depends on the magnetic field B as

$$f_{rev} = \frac{cB}{2\pi R_0 \sqrt{(\frac{m_p c}{\rho q_0})^2 + B^2}}, \tag{1}$$

with c being light speed, $R_0 = 25$ m the machine radius, m_p the mass of the proton, q_0 its charge, and $\rho = 8.23887$ m the bending radius of the Booster main magnet (Fig 2).

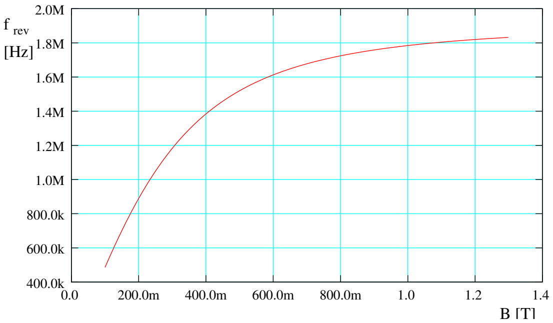

Figure 2: f_{rev} as a function of the field of the bending magnets

A typical magnetic cycle is displayed in Fig. 3. Beam injection takes place at 275 ms (C275) after the start of the cycle, and ejection is at C805.

The peak rate of change of the magnetic field is reached about 600 ms into the cycle, at $\mathrm{d}B/\mathrm{d}t = 2.26$ T/s, and the peak rate of change of the revolution frequency for protons is $\mathrm{d}f_{rev}/\mathrm{d}t = 3.8$ MHz/s, occurring a little earlier. It is understood that these values may need adjustment when the Booster will be upgraded to a 2 GeV peak energy.

3 The current orbit system

Each ring of the machine is equipped with 16 electrostatic pick-ups, each consisting of four interleaved electrodes for the horizontal and vertical displacement

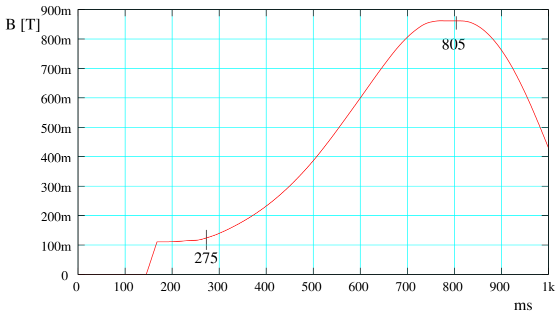

Figure 3: A NORMGPS magnetic cycle

Table 1: PSB parameters at injection

Injection energy P^+:	50 MeV (160 MeV eventually [2])
Injection B-field:	0.1250 T
Injection f_{rev}:	600 kHz
Injection timing:	275 ms

Table 2: PSB parameters for 1.4 GeV ejection

Ejection energy P^+:	1.4 GeV
Ejection B-field:	0.8617 T
Extraction f_{rev}:	1.746 MHz
Ejection timing:	805 ms

Table 3: PSB parameters for 2 GeV ejection

Ejection energy P^+:	2 GeV
Ejection B-field:	1.1275 T
Extraction f_{rev}:	1.829 MHz
Ejection timing:	?

and a dedicated sum electrode. Passive circuitry couples the electrode signals to pairs of 75 Ω coaxial cables that carry the signals up to the surface, distributed in 16 places around the ring in the B37 corridor (Fig 4). That's 24 cables per station. The effective load impedance seen by the electrodes is about 1.4 kΩ, so that the lower cut-off frequency is 500 kHz. At this point, the upper cut-off of about 40 MHz is due to parasitics of the transformers.

Once on the surface, transformers convert the signals to 50 Ω single-ended. The transfer impedance mentioned in Table 4 is defined at this point. The Σ signal level varies from 700 μV_p for the weakest beams, to about 750 mV_p for the most intense. A multiplexer (BRU-MUX) connects the BPM of one ring to a normalizer and provides a copy of the analogue signals to an OASIS crate in BOR (B361 1-0011) [3]. A 40 dB amplifier can be optionally switched into the signal path. This is also where a (sinusoidal) calibration signal can be applied. The two low-frequency normalizer outputs are fed over 50 Ω coax to the BOR into the MPV908 ADCs in Front-End Computer (FEC) dpsbbdr.

Thus, only one of the four rings can be observed at a time. This affects both the orbit measurement and the OASIS analogue signal observation sys-

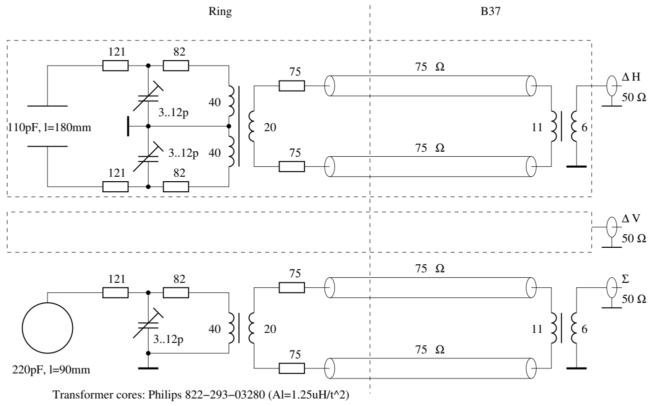

Figure 4: Head electronics schematic diagram. The ΔV channel is identical to ΔH.

Table 4: Booster BPM properties

Electrode aperture	140 mm
Σ electrode length	90 mm
Σ electrode capacitance	220 pF
Δ electrode length	180 mm
Δ electrode capacitance	110 pF
Displacement sensitivity ($\Delta/\Sigma = 1$)	50 mm
Mid-band transfer impedance	170 mΩ (Surface)

tem. A few of the OASIS signals are also used for an experimental digital beam control [4]. Note that this beam control is affected by the settings of the orbit measurement. This interdependency will obviously have to be removed eventually.

Four BPMs, in sections 4, 6, 12 and 14, are used as radial pick-ups. This function must be maintained. In these sections, the BPM signals are daisy-chained through three-way splitters before arriving at the BRU-MUX. Four Δ/Σ demodulators derive a low-frequency radial error signal for each ring on the spot, and these are then fed to the beam-control in the BOR.

4 Basic requirements for a new trajectory measurement

The system will produce turn-by-turn trajectories for individual bunches over the full acceleration cycle. It will also produce time-averaged orbit data for any single bunch and for all bunches together, at a rate of one orbit every millisecond [5]. It will do this all the time, whether or not any clients are requesting measurements. Thus, the system's actions are decoupled from client read-out activities and these clients' actions are also decoupled from one another. Data will be stored in circular buffers large enough to store several cycles worth of data. Thus, data will persist for several seconds before being overwritten with more recent measurements. It is assumed that clients will normally only read subsets of the data that are of interest to them.

A measurement request must specify the following items:

- A cycle sequence or ID number.
- The type of data: Raw, orbit or trajectory.
- The ring number.
- A BPM number.
- A bunch number.
- An event specifier.
- The millisecond timing with respect to the cycle start.
- The number of turns beyond the former.
- The number of data points to return.

The cycle sequence number is assigned by software, e.g., taken from the PLS telegram. Its width may be restricted to just a few bits. It must be wide enough to uniquely identify all cycles for which stored data is available. If the BPM number is different from zero, only data from that BPM will be returned. Bunches are numbered from 1 upwards. The value 0 means all bunches. The exact millisecond timing at which events—e.g., injection or ejection—take place is not very well determined. This problem is solved by allowing the specification of the event starting from which data must be returned.

Depending on the harmonic number and the proportion of ZERO cycles, the overall average accumulation rate of new data is somewhere between 40 and 80 M measurements per cycle, whereas the total average readout rate may only be about 200 k measurements per cycle. Therefore, it will not be possible to read out all data every cycle. It *will* be possible to read out all *orbit* data pertaining to a given cycle, since this will be a data set of small size: 16 BPMs $\times$ 530 measurements $=$ 8480 measurements per ring. It *may* be possible to read turn-by-turn single-bunch position data over a full cycle for one or two BPMs, but this will not be possible for all BPMs, nor will it be possible to do so for several cycles consecutively.

Acquiring all four rings simultaneously affords the possibility of comparing trajectories in different rings within the same cycle. It also provides some measure of decoupling between multiple users wishing to observe different rings. Of course, system cost would be several times that of a single-ring acquisition system.

4.1 Hardware Architecture

The basic architecture for one BPM consists of three high-speed ADCs, one for each of the Σ, ΔH and ΔV signals, an FPGA and a large memory (Fig. 5)[6]. The ADCs deliver a continuous sample stream at a rate of the order of 100 MS/s. The FPGA implements all processing required to derive individual bunch positions, as well as averaged orbit data. This reduces the data rate considerably. The results are stored into circular buffers in the memory, which will be sized to contain several complete cycles of acquired data at any time. Data is stored in the form of Σ, X and Y triplets. The actual calculation of beam positions is deferred to the time the data is read out.

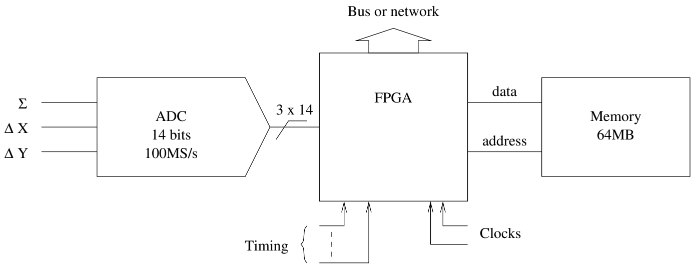

Figure 5: Block diagram for one BPM

Between injection at C275 and ejection at C805, the beam makes about 720 k revolutions. Assuming the machine runs on h=2, that we set aside 16 bits for each of Σ, X and Y, and rounding up to the next higher power of two, each BPM would require 16 MB of data memory per cycle.

Timing and clocks, generated at a central location, will have to be distributed to each channel. The required clocks are the site-wide 10 MHz standard and the machine revolution frequency f_{rev}. The latter needs to be present only for a few milliseconds around injection time. The system relies on the phase of f_{rev} to find the position of the first RF bucket. The required timing events are:

- Start of Cycle
- Start and End of Calibration interval
- Injection trigger(s)
- RF gymnastic trigger(s)
- End of Cycle/Ejection

All timing signals must be aligned to the 10 MHz clock, to ensure that all processing channels act synchronously. This implies the need for a timing synchronizer box that will also serve as a timing distributor. The injection timing is the most critical. It should occur during the revolution period preceding injection, i.e., inside of a window of 1.67 μs. A millisecond clock is not required; It will be generated internally from the Start-of-Cycle and the 10 MHz clock. Note also that no B-train is needed.

4.2 Principle of operation

The system will measure the centre of charge of each bunch. In order to do so, each of the three BPM signals, Σ, ΔH and ΔV are integrated over the length of the bunch. The position is then calculated via

$$x = S\frac{\Delta}{\Sigma} + E, \tag{2}$$

where S is a scaling factor derived from the BPM geometry and the properties of its analogue signal conditioning electronics, and E is an error correction accounting for mechanical mis-alignment and electrical offsets.

The integrals are calculated by the hardware simply by adding together the samples belonging to a bunch. To tell which those are, beam-synchronous timing signals are needed. For each BPM station, a numerical Phase Locked

Loop (PLL) algorithm locks a Numerically Controlled Oscillator (NCO) to the revolution frequency f_{rev} prior to beam injection, and to the BPM's Σ signal afterward. All required beam-synchronous signals are then derived from this PLL. All this is implemented in the FPGA.

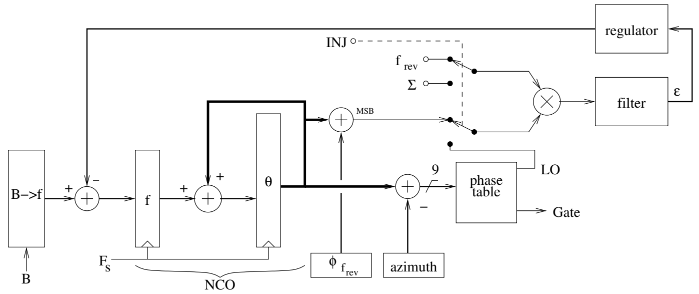

Figure 6: Principle of beam-synchronous signal generation.

Referring to the PLL block–diagram in Fig. 6, a value of f is added to an accumulator register θ each period of the sampling frequency F_s. The initial value f is calculated from the B-field at injection, using relation (1), so that θ overflows at the rate of f_{rev}. A phase table addressed by the value of θ contains h periods of a Local Oscillator (LO) signal, which is mixed with the BPM's Σ and low-pass filtered to derive a phase error ε. A feedback regulator corrects the value of f such that ε is driven towards zero, thereby locking θ to the phase of Σ [6].

Another output of the phase table produces a *Gate* signal that designates the samples belonging to each bunch. At the end of each gate, the accumulated integrals are stored and the integrators are reset.

The phase table is initialized prior to each acceleration cycle with LO and *Gate* signals appropriate for the type of cycle coming up.

The NCOs for all BPM stations run at the same phase. An azimuth value is subtracted from θ to properly phase the beam-synchronous output signals according to the position of the BPM in the ring. Thus the phase tables can be identical for all BPM stations.

Also, timing events are resynchronized with an overflow of the phase table address, with the effect that timing events seem to propagate around the ring in step with the beam, starting from sector (or period) 1. A value $\phi_{f_{rev}}$, common to all BPM stations, adjusts the phase of the NCO with respect to that of the f_{rev} input, so that timing events affect the BPM station with the smallest azimuth first.

4.3 State machine

Each cycle is subdivided into a calibration interval, a beam acquisition interval and several transition intervals (Fig 7). A state machine, driven by accelerator timing events, steps through the successive states assigned to each of these intervals, selecting phase table columns containing the appropriate beam-synchronous signals for each interval. The state machine traverses a programmable sequence determined by a state-transition table, which is basically a matrix which for each state and for each possible event specifies the next state

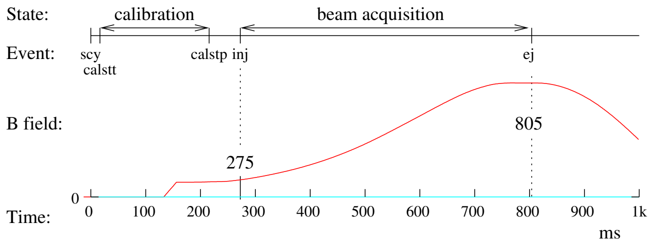

Figure 7: States and events along the cycle

to go to. As explained in the previous section, the state changes propagate around the ring in step with the beam, even though the events reach all BPM stations simultaneously.

4.4 Initialization issues

The system needs information from the accelerator control system in order to select the right mode of operation. The PLS telegram is used to select one of a set of previously built configuration files. The configuration files provide state-transition tables, phase table images and initial NCO frequency values with their respective load times. Software loads the information from the selected file into the system just before the start of a new cycle.

Preparation of the configuration files is the system expert's job, although a tool to allow the control room crew to edit them can be envisaged. In the interest of decoupling different clients, we should resist the temptation to allow operators to change settings affecting multiple clients.

It must be possible to initialize the value of the NCO frequency f to two independently settable values at two independently programmable instants in the cycle. A 1 ms timing resolution is sufficient. One of these values would be the revolution frequency of a simulated beam signal for calibration, and the other would be the expected revolution frequency at actual beam injection. The value of f will be initialized to either of these values a millisecond or two before calibration or injection, leaving the PLL some time to align its phase with that of the f_{rev} signal.

4.5 Data structures

In the course of its operation, the hardware builds and maintains several data structures:

- The trajectory data, its principal product.
- Orbit data, averaged position stored every ms.
- An event index, containing for each timing event the cycle sequence number, state, filling pattern and start address of associated trajectory and orbit data.
- A millisecond index, containing the cycle sequence number and the trajectory data memory address reached at each millisecond tick.

All these are organized as circular buffers. It may be useful, but probably not indispensable, to have the respective sizes of these circular buffers programmable by software. Both indexes have readable pointers, which serve as starting points for the software to access the data. Using the information in these tables, software can locate the data pertaining to a given measurement request, if it still exists.

4.6 Diagnostic data

In order to trouble-shoot, adjust and optimize the system's operation, the ability to acquire a certain amount of diagnostic data is crucial. A number of relevant internal signals must be sampled at a selectable rate from F_s downward in a 1-2-5 or binary sequence. Only about 1 k samples need to be recorded, although more is of course welcome, as FPGA resources allow. The slowest rate needed is such that the recording comfortably spans a full acceleration cycle. The required signals are:

- The input f_{rev} signal. (1 bit)
- The most significant bit of the PLL phase accumulator θ.
- The BPM Σ signal. (8 bits)
- The LO, Gate and BLR signals. (1 bit each)
- The phase error ε. (8 bits)
- The PLL frequency f. (8 bits)
- The current state. (4 bits)

These signals are continuously streamed at the selected clock rate into a fixed-size circular buffer memory. The buffer is frozen after a delay of a programmable number of periods of F_s following the occurrence of a programmable trigger event. The range of the trigger delay counter should be large enough to span a full acceleration cycle. A trigger event may be any of the input machine timing events, or a transition into a given state, either alone or in combination, and possibly conditioned by a mask on the sequence number. The possibility to have the initial frequency load events as diagnostic triggers may be useful too. The data in the buffer is to remain frozen until it is either read out, or until any of the trigger or timing conditions are changed.

5 Calibration

The desired absolute position accuracy is 300 μm, and the resolution 200 μm [5]. An estimate based on experience with existing systems operating with similar beams and pick-ups shows that the resolution requirement can be fulfilled. A more accurate assessment will have to await the detailed hardware design. Due to the effective averaging of trajectories to obtain the orbits, the resolution for orbits would be much better.

The absolute position of the BPMs is not very well known. The installed monitors do not have enough alignment targets or reference surfaces or edges. Therefore, absolute alignment cannot be assured using the usual metrological methods. Removing the BPMs from the ring to be measured on a test bench—which doesn't exist—is deemed to be risky, especially considering that there are no spare BPMs. Can we do something using beam-based alignment?

Electronics calibration can be performed—on demand—in the interval between Start-of-Cycle and beam injection. This interval is 275 ms long, allowing plenty of time. Thus, calibrations need not interfere with normal measurements.

A calibration generator sends beam-like pulses to each of the BPMs in such a way as to simulate a beam in the machine at one of two well-known positions with respect to the BPM electrical centre. This is acquired by the system exactly as would be a real beam. Scaling and offset values are then calculated such that the measured data yields the well-known positions, using standard linear regression. This is done for all possible operational gain settings.

6 Front-end software

The front-end software is responsible for programming the hardware to execute the appropriate sequences on the one hand, and for communicating measurement results to clients on the other hand. The former task is synchronous with the accelerator's operation, while the latter depends principally on client's actions. This suggests a natural distribution of work between one real-time thread and as many server threads as there are outstanding client requests.

The tasks of the software can summarily be broken down as follows:

- From the PLS telegram, determine the required system settings and apply them in the interval between beam ejection of one cycle and start of the next. Gain settings, state-transition tables and phase table data are kept in files mapped to the `USERline` field of the PLS telegram. While state and phase settings must be applied prior to the start of a new cycle, BPM amplifier gain settings must be applied in the interval between the end of calibration and beam injection.
- Accept client's measurement requests, check for sanity and collect and send data when available. Convert raw data to positions by applying scale and offset values on the fly.
- On demand, collect calibration data and calculate and store new scale and offset values. Apply sanity checks. Amplifier gain settings during calibration may be different from those used during beam acquisition and should be applied during the interval between ejection and start of cycle.
- On demand, collect diagnostic data and send it to the client when available.
- Provide access to and allow modification of state-transition and phase tables, gain values, values and timings of injection frequencies and calibration values. This is probably best done via a separate program.

In order to conserve a responsive interface, requests that deliver data which do not depend on accelerator timing events should be delivered asynchronously. (The application level program shouldn't 'hang' while waiting for events or data requests to be fulfilled.) Care should be taken to ensure that the application level programs are kept free from real-time constraints.

7 Front-end hardware

The BPM signal processing hardware must be extensively changed. The passive hybrid circuitry (Fig 4) of the current installation is ageing and needs to be replaced. The 75 Ω coaxial cabling between the ring tunnel and the surface can probably be re-used. Additional amplifiers, with gain settable in roughly 6 dB steps will be needed. These can be installed in the existing NIM crates distributed around the B37 corridor. A new calibration generator is needed too, producing beam-like signals rather than sinusoids. The acquisition hardware is assumed to be based on FMC ADC plug-ins, of which two fit on a VME FMC carrier board [7, 8]. This board also contains the FPGA that will handle the signal processing. Finally, a timing distributor and fan-out box will be required. This device will produce reference frequencies for the acquisition modules and the test generator and it will synchronize external timing signals.

Two possible implementations are proposed for the new installation: The complete four-ring version and the "economy" single-ring option. The former, Fig 8, assumes individual acquisition chains for each of the rings. This obviously means that many hardware components come in fours, and it further implies the installation of quite some extra cabling which ends up dominating the cost (Table 5). Note that this assumes that the RF cabling between the tunnel and the surface will be reused. If the decision is taken to replace this cabling too, this would roughly add 100 kSFr to the RF cabling costs for both single and 4-ring versions.

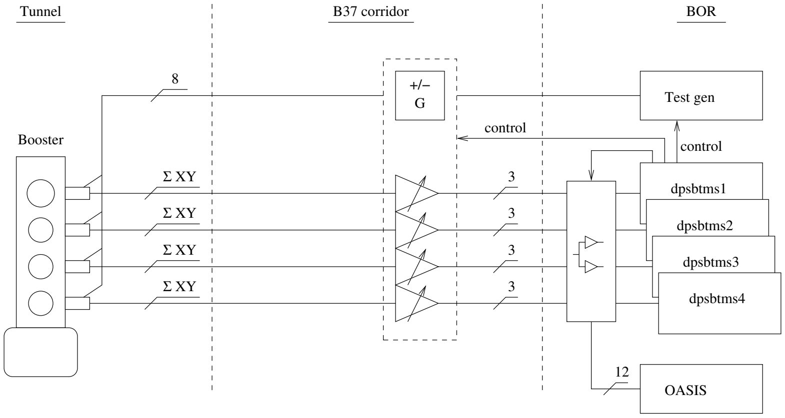

Figure 8: Layout of the full four-ring system for one BPM. Four FECs (dpsbtms1-4) are needed, one per ring.

Distribution amplifiers provide a copy of the 192 analogue signals to OASIS. If the number of OASIS channels cannot be increased beyond the 48 currently installed, some multiplexing will have to be done in the distribution amplifiers.

The single-ring installation, Fig 9, uses one single acquisition chain and re-uses the BRU-MUX module to connect it to the desired ring. In addition, the re-use of existing RF cabling between the B37 corridor and BOR is assumed. Only a few new cables are needed and the cost is now dominated by the front-end passive hybrid circuits (Table 6). The cost figures are, of course, approximate.

Table 5: Cost breakdown for the four-ring version

Device	Unit cost	Qty	Totals
Hybrid	1000	64	64000
RF cabling B37-BOR	400	256	102400
Variable gain amps (3 ch)	1000	64	64000
Gain control interface	1500	4	6000
Gain control cabling	400	16	6400
FEC	10000	4	40000
VME FMC carrier	2200	32	70400
FMC ADC piggy-back	500	64	32000
Test generator	15000	1	15000
Distribution amps (per ch)	100	192	19200
Timing sync + dist	15000	1	15000
Misc, patch cabling, etc.	5000	1	5000
Total			439400

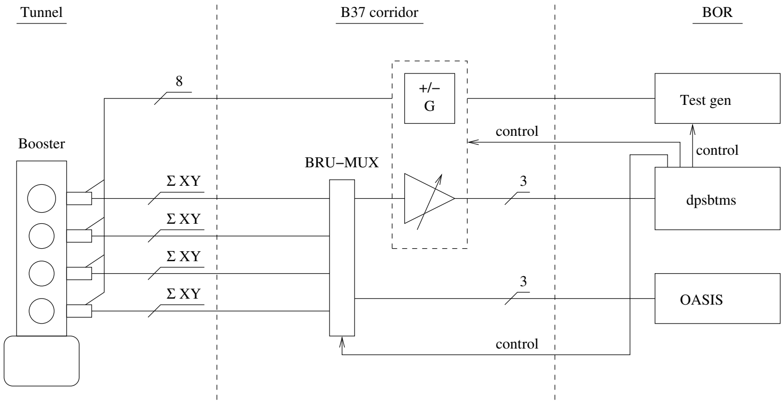

Figure 9: Layout of the single-ring system for one BPM. One FEC (dpsbtms) is needed for the whole system.

Table 6: Cost breakdown for the single-ring version

Device	Unit cost	Qty	Totals
Hybrid	1000	64	64000
RF cabling B37-BOR	400	32	12800
Variable gain amps (3 ch)	1000	16	16000
Gain control interface	1500	1	1500
Gain control cabling	400	16	6400
FEC	10000	1	10000
VME FMC carrier	2200	8	17600
FMC ADC piggy-back	500	16	8000
Test generator	15000	1	15000
Distribution amps (per ch)	100	0	0
Timing sync + dist	15000	1	15000
Misc. Patch cabling, etc.	2500	1	2500
Total			168800

8 Conclusion

This note describes the design for a new Booster trajectory Measurement System, very similar to that of the PS in its principle of operation. Contrary to the PS system however, the implementation proposed here is based on standard BE/CO hardware: FECs and BE/CO-supported acquisition modules. Two possible versions are examined: A single ring acquisition system that can be connected to any one of the Booster rings, and a complete system, in which every Booster ring has its own acquisition chain. Approximate cost for the former is 170 kSFr, and for the latter 440 kSFr.

References

[1] W. Weterings, M. Aiba, J. Borburgh, C. Carli, T. Fowler, B. Goddard, "Operational considerations for the PSB H- Injection System", sLHC-PROJECT-Report-0006, May 2009.

[2] F. Gerigk, C. Carli, R. Garoby, K. Hanke, A. M. Lombardi, R. Maccaferri, S. Maury, C. Rossi, M. Vretenar, "Construction Status of Linac4", sLHC Project Report 0027, May 2009.

[3] F. Guillot-Vignot, "Booster ring pick-up (BRU) multiplexer", CERN-BE-Note-2010-009 BI.

[4] M.E. Angoletta, A. Blas, A. Butterworth, A. Findlay, P.M. Leinonen, J.C. Molendijk, F. Pedersen, J. Sanchez-Quesada, M. Schokker, "CERN's Booster LLRF renovation: Plans and initial beam tests", IPAC10, TU-PEA056.

[5] C. Carli, B. Goddard, K. Hanke, B. Mikulec, "Specifications for PSB ring Pick-ups with Linac4 as injector", `https://edms.cern.ch/document/1078894`.

[6] J. Belleman, "A New Trajectory Measurement System for the CERN Proton Synchrotron", CERN-AB-2005-059.

[7] J. Serrano, P. Alvarez, "VME FMC carrier", `http://www.ohwr.org/projects/fmc-vme-carrier>`.

[8] M. Fimiarz, M. Cattin, T. Wlostowski "FMC ADC 100M 14b 4ch", `http://www.ohwr.org/projects/fmc-adc-100m14b4cha`.