

LCLS Stripline BPM Processor Engineering Specification

Till Straumann

INTRODUCTION

This document specifies the LCLS BPM signal processing system needed in order to meet the physics requirements for stripline BPMs outlined in PRD 1.1-314. Since those requirements call for using existing stripline devices in the LINAC and LTU (FFTB-type striplines), the design of the stripline assemblies falls outside of the scope of this document.

However, here all other components such as cables, conditioning and acquisition electronics, calibration/self-test features, data processing and the control- and timing system interfaces are described.

SYSTEM OVERVIEW

Fig. 1 shows a block diagram of the BPM processor. The raw stripline signals travel through a “long-haul” cable out of the accelerator tunnel, where no active parts are to be located, before reaching the processor. Self-test couplers allow for the injection of test signals upstream of all electronics in order to verify and calibrate their operation. The analog front-end conditions the signal for digitization by selective filtering and amplification. A variable gain stage which is remote-controllable via the control system can be used to optimally adjust the signal level for the ADC. Finally, a commercial wide-band ADC is used to simultaneously digitize the signals from all four strips and making the digital values available to a CPU.

The CPU implements the necessary low-level data processing the PRD 1.1-314 calls for (such as timestamped ring-buffers, calibration, averaging etc.) as well as an EPICS IOC to make all the functionality available to the control system.

In order to meet the real-time performance needed for triggering, buffering and timely data acquisition, a hard real-time operating system is employed which schedules the critical tasks in a deterministic manner.

A “event-receiver” module is integrated into the processor providing the ADC gating, timestamp information and interrupts for task scheduling etc.

An auxiliary fast communication port is available for propagating BPM data to other subsystems such as feedback applications without incurring the overhead and non-deterministic latencies of the networking stack and physical media.

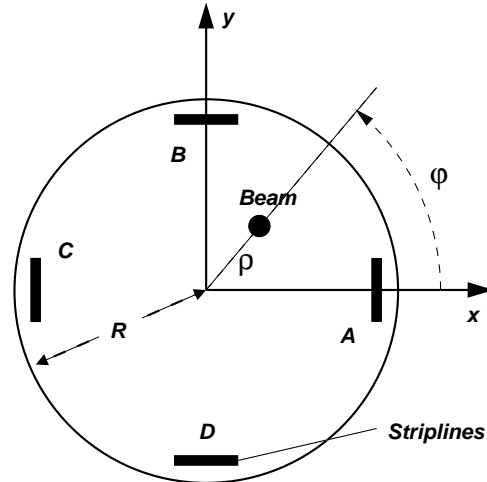


Figure 2: BPM transversal geometry

SYSTEM DESCRIPTION

The central functionality of the BPM processor is determining beam position from stripline signal data. Therefore, the properties of those signals shall be reviewed.

Stripline Signals

As a result of the very short LCLS beam passing by a stripline detector, a pulse doublet is induced at the upstream end of the long-haul cable. Over a bandwidth exceeding the range of interest (>1 GHz) by far, the spectral distribution of the stripline signal (Volts/Hz) can be described by

$$Q Z F \left(1 + 2 \frac{\rho}{R} \cos(\phi) \right)$$

with

- Q beam charge
- Z line impedance (50Ohm)
- R beam pipe radius
- β half strip angular coverage ($\frac{\omega}{2} R$)
- l strip length
- c speed of light
- ρ beam displacement from center
- ϕ beam azimuth against strip center (fig. 2)

and the “form-factor” F

$$F = \frac{\beta}{\pi} \sin(\omega l / c)$$

Its frequency dependence has to be taken into account when choosing an appropriate processing frequency.

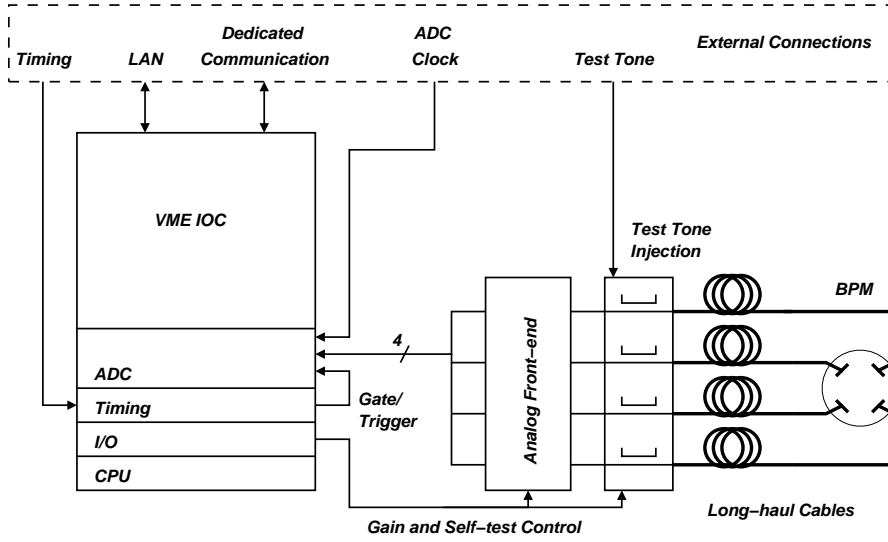


Figure 1: BPM processor block diagram

Resolution

The sensitivity of the detector with respect to positional changes can be obtained by normalizing the difference of the signal amplitudes A , C of two opposite strips ($\phi_C = \phi_A + \pi$) to their sum:

$$\frac{A - C}{A + C} = 2 \frac{\rho}{R} \cos \phi_A \quad (1)$$

Digitizer Resolution The ratio eq. 1 also translates the physical resolution requirement into a resolution requirement for the BPM processor. In the LINAC, where $10\mu\text{m}$ are required and where the pipe diameter is $\approx 1''$, a resolution of better than $8 \cdot 10^{-4}$ which corresponds to 10.3 bits is required. One additional bit is required to accommodate for the required excursions to $\pm \frac{1}{3}$ of the stay-clear area. A state-of-the art 100MSPS digitizer provides only ≈ 11.3 effective bits. If an expected processing gain of 1.5 bits (corresponding to oversampling in the order of a factor of 10) is added it becomes obvious that such a device just barely can meet the requirements.

The variations in beam charge over a range from .2..1nC the PRD 1.1-314 specifies can only be handled by electronically switching the front-end gain or by employing a hybrid to provide separate “difference” and “sum” signals. It is assumed that no frequent or sudden changes of beam charge occur or that the resolution requirements are relaxed when such changes are expected.

Analog Frontend Resolution One purpose of the analog frontend is reducing the signal bandwidth to a level that can be handled by the ADC. Reducing bandwidth significantly below the Nyquist limit would help improve the digitizer resolution due to increased processing gain. However, as a statistical analysis shows, reducing the system bandwidth B_a also reduces the SNR of the analog signal according to

$$\frac{1}{2} \frac{\delta \rho}{R} \approx \frac{1}{2\sqrt{2}} \sqrt{\frac{kT N_F}{Z (FQ)^2 4 B_a}} \quad (2)$$

where the SNR has been expressed as a statistical variation of apparent position (normalized to the radius R). The factor of $1/2$ on the left side is introduced by the requirement that the noise floor RMS shall not exceed half of the positional resolution. N_F denotes the total noise figure of the frontend, i.e., including cable losses.

Design of the analog frontend has to be done carefully in order to preserve linearity and a low noise floor as the signal passes through several stages of filtering and amplification. (As the bandwidth of the initially very short, pulsed signal is reduced, its amplitude decreases. Neither must the first stages go into saturation nor must they increase the noise floor).

Long-haul Cables

Losses in the long-haul cables increase with frequency. On one hand, this is desirable since it reduces the high frequency content ($< 1\text{GHz}$) and hence the load on the frontend. On the other hand, valuable signal power is lost, too, which results in a reduction of SNR and hence resolution. Cable losses have to be considered when choosing the processing frequency.

Self-Test/Calibration Feature

Remotely controlled, a CW signal generator can inject a tone individually, at each of the four strips. Picking this tone up from the three other striplines can provide important diagnostic information.

- test functionality of strips, connectors, cables etc. down to the ADCs.

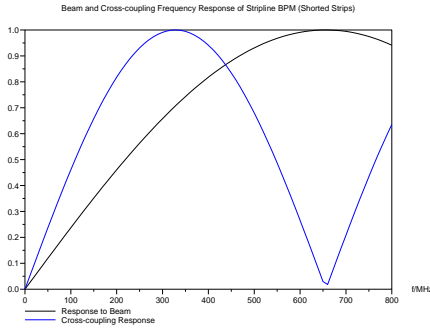


Figure 3: Frequency response to beam and calibration tone

- gain calibration. To the extent the stripline arrangement is symmetrical with respect to the mechanical center, absolute offset calibration can be performed, i.e., the relative channel gains can be adjusted to remove any apparent offset.

Even if symmetry is not perfect, as long as the properties of the assembly are constant, it is still possible to compensate for drift of the relative gains of the electronics.

Since the duty cycle of the BPM processor is very low (ADC gate vs. 8ms repetition rate) there is plenty of time available for scheduling calibration/tests between beam pulses.

It should be noted that the frequency response of the cross-coupling between strips is quite different from the response to an actual beam. Furthermore, it depends on the stripline termination which is assumed to be a short-circuit to the vacuum pipe. In this case, the cross-coupling exhibits a

$$\sin(2\omega l/c)$$

behavior, i.e., at the frequency of maximal response to the beam there is *zero* (see fig. 3).

SYSTEM DESIGN

After introducing the core system components, we are in the position to summarize the key design parameters.

Analog Components

Cables Cable losses vs. frequency and the stripline frequency response weighted with cable losses are depicted in fig. 4 for several cable types (including the RG-223 currently installed in the LINAC). Cable data will be needed when determining the analog bandwidth and processing frequency.

Analog System Bandwidth Eq. 2 is employed to calculate the resolution that can be achieved for a given stripline geometry, beam charge, noise figure, bandwidth and processing frequency. The positional uncertainty (note the PRD 1.1-314 calls for this number being less than half

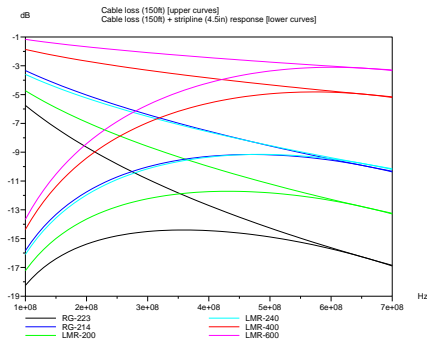


Figure 4: Losses of different types of cable. Stripline frequency response weighted with cable losses are also shown.

the required resolution) has been tabulated for a number of parameters in tbl. 2.

Obviously, it is not trivial to meet the requirements at low bunch charge. With currently available digitizers it is not possible to increase B_a beyond $\approx 5 - 10$ MHz (some processing gain is needed to meet the resolution requirements on the digital side, see above). Hence, the system noise figure must be minimized by choosing good cables and a low-noise design of the front-end electronics. Also, a higher processing frequency may be chosen (up to ≈ 400 MHz) but above ≈ 180 MHz a downconverter must be implemented and the required LO drive signal needs to be generated and distributed. In addition to operating at a frequency where more signal power is available, a downconverter would have the additional advantage of converting to a IF in the first Nyquist zone where the ADC performs better.

Note (see fig. 4) that in case of the legacy RG-223 cable, increasing f_0 does not help due to the high losses.

Because it might be desirable to use different f_0 for different BPMs (e.g., the FFTP-style BPMs are much longer), the analog front-end should be designed in a way that allows selection of f_0 by exchanging a single filter. This should not be difficult to achieve due to the availability of cheap broad-band components in the frequency range of interest.

Digitizer

Sampling Frequency Obviously, sampling frequency should be as high as possible. As will be shown below, at low charges, the resolution requirement implies a certain signal bandwidth B_a of multiple MHz. Such a signal can only be digitized by modern high-speed devices.

ADCs ranging from $ENOB \approx 7 @ 1000$ MSPS, $ENOB \approx 10 @ 200$ MSPS and $ENOB \approx 11 @ 100$ MSPS are widely available today (i.e., “boards” not just “chips”). Oversampling can be employed to increase the number of effective bits. However, a factor of 4 is needed to gain an additional bit. Therefore, a 14-bit digitizer with $ENOB \approx 11$ and a sampling rate of

Sampling Frequency	f_s	Chose $ENOB + \frac{1}{2}(\log_2(f_s/B_a) - 1)$ as high as possible. 4-times oversampling is needed to gain one $ENOB$.
Processing Frequency	f_0	A higher frequency yields more signal but emphasizes cable losses. If f_0 stays below the input bandwidth of the ADC a mixing stage can be avoided (undersampling). f_0 should not be close to the maximal response because the test-tone response vanishes. $f_0 \pm B_a/2$ must also stay away from multiples of the Nyquist frequency to avoid undesired aliasing.
Processing Bandwidth	B_a	High bandwidth increases (analog) SNR and resolution. Low bandwidth increases digital processing gain (but this cannot “restore” lost analog SNR).
Noise Figure	N_F	Should be as low as reasonably achievable. Contributors are cables, first stages in front end.
Number of Bits	$ENOB$	Effective number of bits of the ADC. A higher f_0 and/or a poor clock usually results in a reduction of $ENOB$.

Table 1: Important System Parameters

Stripline	f_0	cable	B_a	$\delta\rho$
Injector ($l = 3.8''$, $D = 1.37''$)	140MHz	200' lmr400	5MHz	14 μ m
Injector ($l = 3.8''$, $D = 1.37''$)	140MHz	50' lmr400	5MHz	11 μ m
Injector ($l = 3.8''$, $D = 1.37''$)	140MHz	200' lmr400	15MHz	8 μ m
Injector ($l = 3.8''$, $D = 1.37''$)	300MHz	200' lmr400	5MHz	8 μ m
LINAC ($l = 4.5''$, $D = 1.05''$)	140MHz	150' rg223	5MHz	15 μ m
LINAC ($l = 4.5''$, $D = 1.05''$)	140MHz	150' rg223	15MHz	9 μ m
LINAC ($l = 4.5''$, $D = 1.05''$)	140MHz	150' lmr400	5MHz	9 μ m
LINAC ($l = 4.5''$, $D = 1.05''$)	140MHz	150' lmr400	15MHz	5 μ m
LINAC ($l = 4.5''$, $D = 1.05''$)	300MHz	150' lmr400	5MHz	5 μ m
LTU (FFTB) ($l = 21''$, $D = 0.786''$)	70MHz	150' lmr400	5MHz	3.9 μ m
LTU (FFTB) ($l = 21''$, $D = 0.786''$)	70MHz	150' lmr400	15MHz	2.2 μ m

Table 2: Positional uncertainty due to limited SNR. (For the injector striplines, angular coverage comparable to the LINAC striplines (7%) has been assumed; in case of the FFTB striplines a number of 5% was used). System impedance is 50Ohm, the numbers were calculated for a noise figure of 10dB (translates directly into position; reduction of N_F by 3dB yields an improvement by $\frac{1}{\sqrt{2}}$).

105..125MSPS seems to be the best choice.

The sampling clock has to be of low jitter in order not to further reduce the *ENOB*. Phase noise in the frequency band $B_a/10..f_s/2$ (and its aliases) is of concern, not so much the close-in phase noise which is more difficult to reduce.

Gating

A programmable “event-receiver” module connected to the LCLS timing system provides the necessary functionality to synchronize gating for the ADC with beam-arrival.

- Exact timing and width of the gate needs to be software controllable.
- Software needs to be able to perform data acquisition independently of the main gate in order to implement the calibration/self-test features.
- It needs to be verified that event-receiver output signals are compatible with the ADC module (logical levels etc.).

Finally, it should be noted that the PRD 1.1-314 specifies an unreasonable maximum on the gate width of “not significantly more than 8.4ns”. Obviously, the gate width must be $>\approx 5..10/B_a$ for the ADC to capture the full time response and with a limit of $10MHz$ to B_a having been identified as a result of today’s ADC’s limitations, the gate width cannot be less than $\approx 1\mu s$.

CONTROL SYSTEM INTERFACE

The IOC crate controller CPU is chosen to provide enough computing power and memory to provide the functionality required by PRD 1.1-314. The 120Hz repetition rate is not very challenging for modern hardware and a real-time operating system. The system should be able to provide BPM readings with a latency of $< 300\mu s$ with respect to beam arrival. Note that this figure is largely dominated by VME backplane transfers. If multiple BPMs are processed in one crate, in the order of $200\mu s$ of transfer time have to be budgeted per BPM.

If other VME modules are present (such as the timing module), drivers must be carefully written and VME access appropriately scheduled to avoid stalls that could further increase latency.

Some vendors offer ADC cards with Race++ interface which significantly boosts transfer speed. Race++ requires a special, additional backplane to be installed in a VME64x crate and a PMC adapter on the CPU card.

Real-time tasks with appropriate priorities take care of calculating positions, calibration/self-test, maintaining ring-buffers, running averages, timestamping data etc. and sending out information using the dedicated communication port.

At a higher level, EPICS IOC software is responsible for I/O to the user. All control parameters and monitored

data of the real-time controls are accessible via EPICS and the channel access protocol, i.e., except for the dedicated real-time communication channel no non-standard protocols will be used.

A middle-layer of software integrates the real-time data acquisition with the “SLC-aware” portions of the IOC.

Since this is a real-time critical application, it is required that no software components to be loaded on the BPM processor (including third-party drivers, timing software, other IOC components etc.) execute any code from ISR context except for task synchronization primitives. This policy ensures that all latency requirements can be met by means of assigning appropriate scheduling priorities.

The PRD 1.1-314 requires some “user-specific” settings to be maintained at the IOC level. Since in general, resources on an EPICS IOC are never tied to a specific user and the whole concept of a “user-specific” resource is not built-in this requirement cannot be met (except perhaps in the SLC-aware portion). However, it is possible to provide multiple instances of critical resources (such as “freezable ring buffers”) but access policies to any EPICS database items are beyond the scope of the BPM processor.

REQUIRED BPM PROCESSOR INSTANCES

A single BPM Processor requires at least the following hardware

- Remote-controllable VME crate with CPU.
- Low-end digital (and TBD analog) I/O card or IP module for control and monitor signals.
- Event receiver card to interface with the timing system.
- High-speed ADC card. Commercially available cards usually have 8 channels, enough for 2 BPMs. Additional BPMs can be processed in the same crate up to a limit that is set by latency requirements.
- High speed LAN access; depending on usage 100Mb/1000Mb.
- Terminal-server with access to CPU diagnostic port.
- Analog front-end (4 channels per BPM), including power supply. (Form factor TBD.)
- Calibration distribution and injection hardware (TBD).
- ADC clock distribution. Alternatively, a suitable VME or PMC module could be used.

The total count of BPM processors (TBD needs to be double-checked) is as follows

- Injector: 25 BPMs; depending on tolerable latency, these have to be split into 3-6 processors. In the injector, where high density is an issue it might be worthwhile to consider Race++.
- LINAC: There are processors per sector with 8 BPM per sector. I.e., each processor will have to handle 2-3 BPMs. In total, there are 33-34 processors required to instrument the LINAC.
- LTU: There are 3 locations available to instrument 29 BPMs. Handling 10 BPMs per processor should be possible if the associated latency is admissible (≤ 2 ms). Otherwise, using more processors or the Race++ option needs to be considered.
- Dump: 2 BPMs, 1 processor.

In total, less than 41 BPM processors are needed (if a latency ≤ 2 ms can be tolerated).