

The Libera as a PS orbit measurement system building block

J.M. Belleman;

CERN, CH-1211 Genève 23 - Switzerland

T. Karcnik

Instrumentation Technologies, SI-5250 Solkan, Slovenia

Abstract

During the year 2004, extensive tests using a Libera data processor have been made in order to study its suitability as a building block for a complete PS trajectory and orbit measurement system. The Libera consists of four fast 12-bit ADCs, a Virtex II Pro FPGA and a large memory. This note presents some of the results of the analysis of acquisitions made on a position pick-up in the CERN PS.

1 Introduction

In the context of a collaboration contract between Instrumentation Technologies, Solkan, Slovenia and CERN, a Libera processor has been used to acquire and treat data from a PS pick-up in order to assess its applicability to the construction of a new trajectory and orbit measurement system [1]. The architecture of the Libera processor should allow the new system to offer a considerably enriched set of features as compared to the old.

2 Hardware

Libera is an all-in-one solution for accurate beam position monitoring, and local and global feedback building. A single 1U high 19" enclosure contains analogue and digital boards with a dedicated power supply. These are supported through a linux based housekeeping single board computer. The Libera consists of three basic parts: an analogue board, a digital board and a single board computer.

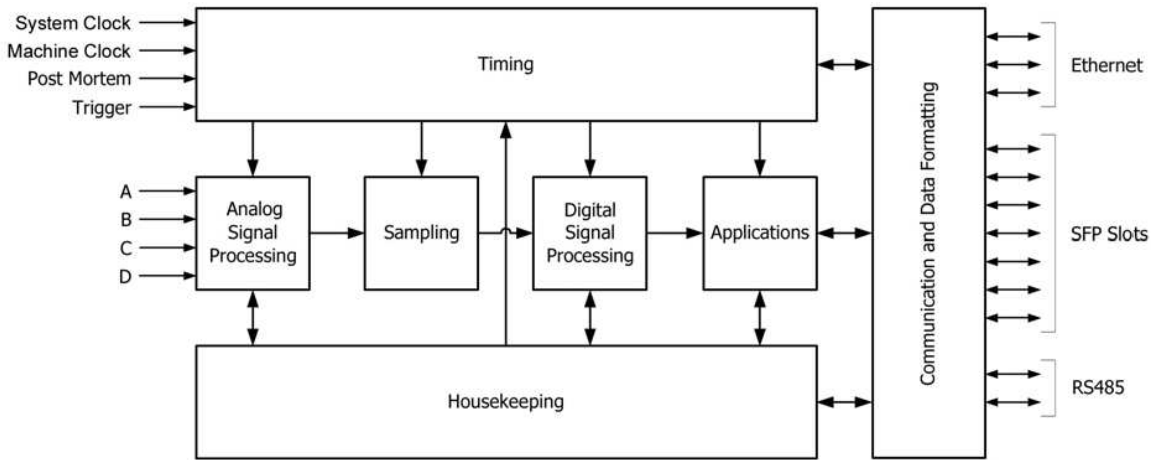


Fig. 1: Libera functional diagram

Fig. 1 presents the Libera functional diagram and Fig. 2 shows its interfaces to the external world. Raw pick-up signals A, B, C and D enter the Libera and are processed, stored and formatted along the way before being sent to the control (Ethernet interface) and/or feedback system (SFP slots). The processing consists of analogue conditioning, A to D conversion, digital signal processing and position calculation. There are two glue functional blocks: timing and housekeeping. Timing interfaces with the external world and in this way provides synchronisation with external events. It is also used for internal synchronisation between different functional blocks and time-stamping of measurements. Housekeeping is a set of specific applications that are, to a great extent, transparent to the control or feedback system, but facilitate performance optimisation and assure continuous and reliable operation of the Libera.

The analogue board is built on a six layer printed circuit board. It has 4 identical analogue inputs and an interface to the digital board. The ADCs (AD9433, 12 bit, up to 125MS/s) sample the pick-up signals at 108.37MHz. The ADC outputs are connected in parallel to the Virtex II Pro chip on the digital board. The inputs are AC coupled. The lower cut-off frequency is below 100kHz and high cut-off frequency is above 50MHz. A 6dB 50Ω attenuator is placed at the input in order to match levels and impedances. The sampling rate is defined by the VCXO-based low jitter (~1ps) oscillator, which can be PLL/FLL locked to the external reference delivered to the "Machine clock" input. Typical tuning range of the VCXO is +/- 50ppm.

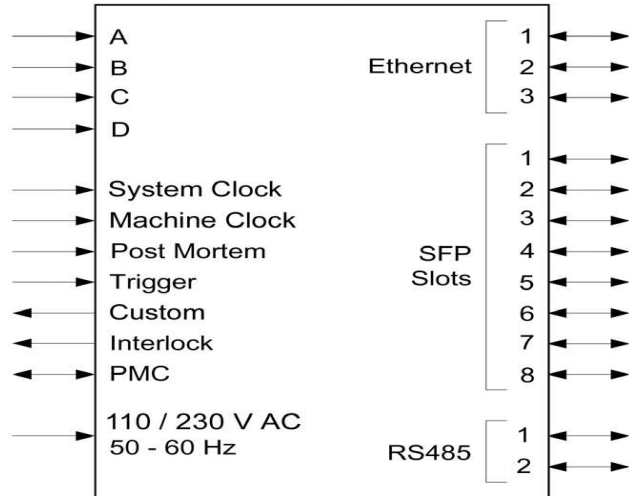


Fig. 2: Libera interfaces

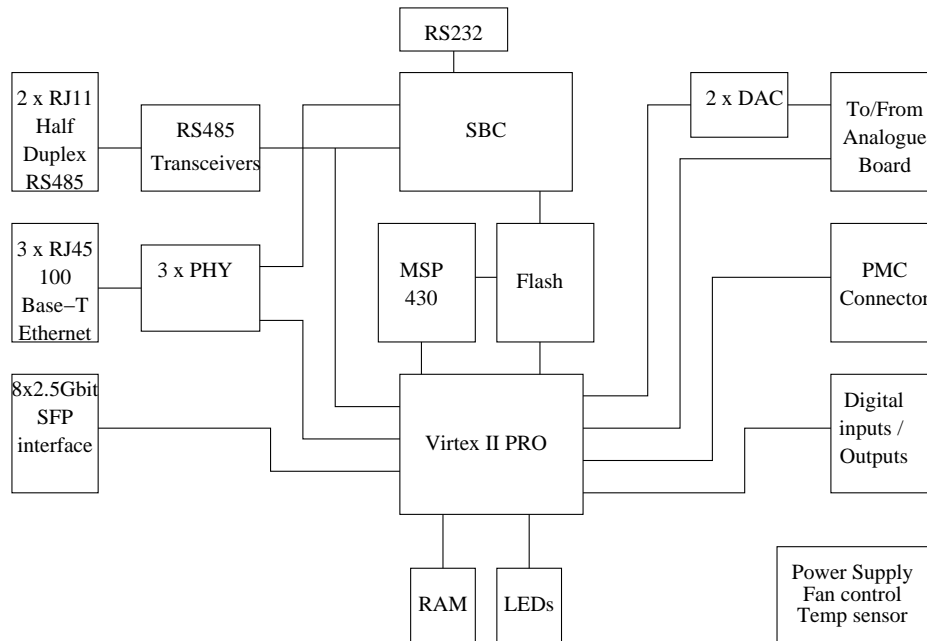


Fig. 3: Digital board block diagram

Fig. 3 shows the Libera's digital board scheme. Core building blocks of the digital board, built on a 14 layer printed circuit, are the FPGA (Field Programmable Gate Array), memory (2 banks of 64MB SDRAM) and a very rich set of interfaces. The Virtex II Pro FPGA with two IBM PowerPC" 405 processors from Xilinx is at the centre of the digital

board. It is a system-on-chip (SoC) solution, which facilitates the following functionalities: digital signal processing, communication, formatting, timing, and part of the housekeeping. This architecture allows replacement of some hard-wired DSP functionality with more flexible procedural description of algorithms. Another powerful feature of the Virtex-II Pro FPGA is integrated up to 3.125 Gbps Rocket IO transceivers. The transceivers address all existing connectivity requirements as well as emerging high-speed interface standards. The Libera also monitors the following components and parameters through I2C:

- Power supply (voltage levels)
- Cooling fans (rpm)
- Libera temperature (digital printed circuit , FPGA, analogue printed circuit)
- Air temperature.

The third component, a single board computer that is built around a StrongArm-based Intel Xscale PXA255A processor, is a mezzanine board. It includes 32MB flash and 64MB SDRAM. It provides application software, bootstrap for the configurable FPGA logic, configuration, diagnostics and maintenance of Libera. The embedded Linux operating system with networking capabilities simplifies integration into an accelerator control system. Fast Ethernet is a native networking solution.

The software is composed of system software and application software. The application software is limited to the SBC. System software, on the other hand, spans the FPGA with embedded PowerPC processors, in addition to the SBC. System software can be roughly divided into an operating system part, covering GNU/Linux and an embedded digital signal processing (DSP) part on the FPGA. Real-time DSP is performed with highly time and space optimised algorithms within the FPGA. For this prototype, this basically means accepting the data from ADCs, storing them in dedicated SDRAM and making them available, on demand, to the SBC software. GNU/Linux operating system on the SBC provides a flexible and capable platform for tasks such as networking, maintenance and customisations. The GNU/Linux package for SBC is based on a standard Linux kernel and includes a file system hierarchy standard (FHS) compliant file system and many common UNIX utilities.

In this prototype the digital board was accessed through the /dev/mem interface, i.e., the digital board was directly mapped to the SBC address space. FPGA registers and functionality are accessed with a user-space “deb” utility. This utility is used to configure the FPGA embedded software, read the acquired data and store it to remote NFS servers for further analysis. A set of additional utilities are available for FPGA reprogramming, etc. Since the SBC software is standard linux based, adaptations are straight forward.

For this prototype, the SDRAM bandwidth was not sufficient to store digitised data from 3 channels at full sampling speed. Thus only two channels were recorded. The new

generation, using DDR2 ram, will solve these problems. The digital board already accepts DDR2 memory.

3 The Proton Synchrotron

The CERN PS is a 200m diameter synchrotron that accelerates protons from 1.4GeV to 26GeV. The revolution frequency of the beam varies from 437 to 477kHz over the course of an acceleration cycle. The protons are injected in bunches and there can be from 1 to 16 bunches injected in one cycle. The PS can produce many beams with differing properties. Each beam type is referred to by its 'user name'. (E.g., EASTB, LHC, etc.)

The PS is equipped with 40 electrostatic position pick-ups (PUs), one of which has been connected to the Libera processor. Each PU has a variable gain amplifier that brings the beam-induced signal up to about 1Vp. The analogue signal bandwidth of the current system is 100kHz-30MHz. The high-frequency roll-off characteristic is Bessel O(5), which is only about 15dB down at 54MHz, the Nyquist frequency at the sampling rate used for this test.

4 Trajectory measurement

The specification calls for the measurement of trajectories, orbits and mean radial position, sometimes per bunch, and sometimes averaged over all bunches [2]. The acquisition of the complete full-cycle trajectory of all bunches, complemented with some post-processing, covers all needs. Assuming that the beam can stay in the machine for 2s, that for each passage of a bunch through a PU, three values are stored (integrals of Σ , Δ_x , Δ_y), that the mean revolution frequency is about 450kHz and that $h=21$, the required memory space per PU is at most 128MB.

For the tests described below, we chose to store raw data. Further processing was done off-line in order to facilitate the study of different algorithms. In the final version, the incoming data would be treated on the fly. In the raw data mode, the Libera, in its present prototype state, can acquire two channels simultaneously to the full extent of its memory (64MB). Since every sample takes two bytes, there are 2^{24} samples per channel. At a sample rate of 108.37MS/s, that's just under 155ms of real time. On an ordinary acceleration cycle, this covers about 70000 revolution periods. One channel was connected to the sum (Σ) signal, and the other to the horizontal difference (Δ).

Fig. 4 shows the raw data taken on an EASTB cycle, with a single bunch at about $5 \cdot 10^{11} p^+$. The PU gain is not optimal for the Libera and the peak excursion of the sum is only about 500 out of a full scale range of 4096 (12-bit ADCs). Nevertheless, the signal is limited by noise that comes from the analogue domain.

In order to find the beam position, both the Σ and Δ signals are integrated over the length of a bunch. The mean position of the bunch is then proportional to the quotient Δ/Σ . To get an accurate integral, it is useful to first restore the baseline. The process is shown in the next few plots. Fig. 5 demonstrates the effect of baseline restoration. Fig. 6




Fig. 4: Raw Σ (red) and Δ (green) signals on EASTB at injection

shows the integrals of the Σ and Δ signals over the first 70171 turns of the beam. The integral was calculated by summing the values of a fixed number of baseline-corrected samples around each bunch. Details of the algorithm are not yet settled and will be the subject of a later publication.




Fig. 5: After application of Base Line Restoration




Fig. 6: After integration




Fig. 7: Bunch positions in mm

Fig. 7 shows the calculated position scaled to mm. The discontinuity at 67500 turns is where the beam goes through transition. The 1σ position resolution is 0.3mm, which is rather poor. The noise spectrum is white. Note that it is the analogue signal that limits the resolution. The Libera itself has already shown its ability to reach the desired 0.1mm resolution [3].

Near transition, the bunch length is considerably shortened as is shown in Fig. 8. The shortest bunch is covered by only 5 samples. Clearly a somewhat higher sampling rate would be beneficial. The overshoot on the short bunch is caused by poor matching of the Libera analogue inputs. This should be easy to fix.




Fig. 8: A bunch near injection (red) and near transition (blue)

Another case that supports the argument for a higher sampling rate is the LHC beam, which runs with 18 bunches at $h=21$ for part of the cycle. Below are plots of a few revolutions (Fig. 9) and a zoom of a few bunches, showing the sample points (Fig.10).

There are only four samples on each bunch. Note also that the reflection following each bunch make the definition of the baseline uncertain.




Fig. 9: LHC beam at $h=21$




Fig.10: Zoom in on LHC bunches

5 Other measurements

The quality of the acquisitions is sufficient for the derivation of other parameters. One example is the machine tune, or Q value. A zoom into the first few hundred turns of an EASTB cycle with one bunch gives a good view of the injection oscillations (Fig. 11). Betatron oscillations provoked by the Q-kicker can also be resolved. The new system can therefore replace the current Q-measurement too. The fact that data from 40 PUs can be combined should lend it an excellent resolution.




Fig. 11: Injection oscillations in the first 200 turns




Fig. 12: Spectrum of injection oscillations

A spectral analysis of these injection oscillations yields a value for the betatron tune of the machine $q=0.16$ (Fig. 12).

6 Beam-synchronous timing

To be able to acquire and store bunch-by-bunch acquisitions, a synchronisation signal that tells where the bunches are is needed. Indeed, given that RF buckets may actually not

be filled, and that the configuration of filled and empty buckets may change without notice, it must be possible to tell where an RF bucket is, even when it's empty³.

The accelerating RF system cannot provide the needed timing because its phase with respect to the beam varies, depending on what is being done with the beam. (Acceleration, coasting, below or above transition.) So except for a brief time prior to injection, the beam-synchronous timing must be derived from the beam itself.

In order to limit the number of potential failure points, one central device should derive the beam-synchronous timing, which would then be distributed to all PU stations. Each station then applies a fixed phase offset appropriate for the PU to which it is connected.

The basic idea of the reconstruction of a phase reference is that of a phase locked loop (PLL) (Fig. 13). Everything can be implemented in digital logic in an FPGA. The oscillator is implemented as a direct digital synthesizer (DDS), but without the sine conversion table. The phase (θ) of the DDS is compared with the PU signal, and a suitably filtered phase error (ϵ) is fed back to keep the DDS synchronized. An initial guess of the revolution frequency (f) is calculated from the measured magnetic field (B) of the accelerator, to make sure that the loop locks to the the right component of the beam spectrum.

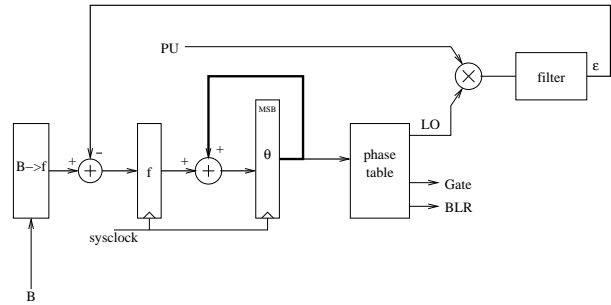


Fig. 13: Principle of reference timing reconstruction




Fig. 14: F_{RF} vs. sample number

Using the acquired data and off-line processing, this method has been put to the test. In Fig. 14, data from two EASTB acquisitions stitched together are used to show that the loop correctly reconstructs the bucket frequency all through the acceleration process.

³ Naturally, at least one RF bucket must be filled, or there will be nothing to be measured.

In Fig. 15, data from an LHC acquisition (six bunches, $h=7$) is superimposed on the reconstructed RF to demonstrate its alignment. The reconstructed RF is in quadrature with the beam because the PLL forces the mean of the product of beam and reconstructed RF to zero.




Fig. 15: PU signal and reconstructed RF at $h=7$

7 Injection and harmonic changes

Prior to arrival of the beam in the machine, there is no PU signal for the PLL to lock on. A signal at the expected revolution frequency, with a known phase with respect to the beam, is substituted as a phase reference. Once the beam is injected, the PU signal can be used.

A number of beams undergo splitting or compression [4][5]. Both operations imply that the harmonic number of the machine is changed on the fly and the system should not lose its synchronism. For example, the LHC beam shortly after injection has six bunches on $h=7$, and later in the cycle, it has 18 bunches on $h=21$. (See Fig. 15 and Fig. 9)

Injection and harmonic change operations are similar in the sense that the system must switch from one reference source or frequency to another, without losing phase lock. To smoothly switch over, the local oscillator (LO) generation, mixer and filter stages are duplicated (Fig. 16). While one is active keeping the loop locked, the other is programmed in preparation for the new conditions. A pulse from the timing system then effects the actual switch-over at the appropriate instant. In the actual implementation, this can be absorbed into the LO generation.

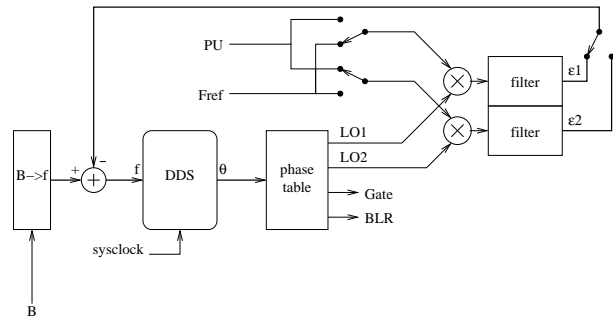


Fig. 16: Coping with RF gymnastics

Returning to the LHC example, Fig. 17 shows a plot of the evolution of the magnitude of spectral lines at $h=7$, $h=14$ and $h=21$ across the bunch splitting operation. An external trigger




Fig. 17: Evolution of magnitude of harmonics in the beam signal for LHC

in the vicinity of the point where the $h=21$ line exceeds the $h=7$ line will effect the change of LO frequency. In this example, there remains quite some energy at $h=7$ after splitting the beam. This strongly depends on the detailed beam structure and cannot be relied upon.

8 Conclusions

We have explored the possibility of using a Libera processor to acquire pick-up signals in the CERN PS. The goal was to have the Libera digitise three analogue channels with 12 bit resolution at a rate of 125Msamples/s and store up to 64MB of data. An external trigger was to start the process. In the end, due to bandwidth problems to the main memory, we settled for a sampling rate of 108.37Msamples/s and two channels. This is a minor issue, as the final system architecture only calls for the storage of FPGA-processed data from three channels at an aggregate rate of 30MS/s, well within the Libera's reach.

Data analysis is still going on. It is clear that the 12 bit resolution of the Libera is amply sufficient. The signal quality is limited by analogue signal noise. The target of a single-pass position resolution of 0.1mm can be attained for most beam types.

Even though a 125MS/s sampling rate is sufficient most of the time, a somewhat higher frequency would be welcome. For the shortest particle bunches observed, e.g., near transition, each bunch is covered by only four or five samples. An RF period at $h=21$ and at 26GeV is only about 11 samples long. The analogue signal bandwidth, about 35MHz and Bessel filtered, is less than a factor of two below the Nyquist limit.

The fact that the sampling frequency is fixed, and not coherent with the machine RF, causes the calculation of the centre of charge of each bunch to become rather noisy. However, the same noise affects the difference signals in the same proportion, and the effective position resolution does not appear to suffer.

The full system will comprise forty LiberAs. Each will acquire the data of a single PU. In order to display meaningful trajectories, data must be collected from each of the LiberAs and assembled into a single measurement. An extra Libera, or more likely a PS standard DSC, can serve that function. It is not clear if Libera's Xscale processor has the necessary throughput for this functionality. The FPGA contains a pair of PowerPC cores that could be put to this task. However, none of this has been tested.

Another issue is beam-synchronous timing. All LiberAs will need to have a common timing reference, not only in terms of ADC and FPGA clocks, but also in terms of accelerator RF phase and C-timing. The details of how this must be done remain to be determined, but preliminary trials with real acquired PU data and off-line analysis indicate that this problem can be solved satisfactorily.

References

- [1] "Design study for a trajectory and orbit measurement system for the CERN Proton Synchrotron", Contract No. K1053/AB
- [2] E. Bravin, J. Belleman, M. Chanel, M. Ludwig, E. Metral, G. Metral, J-P. Potier, U. Raich, R. Scrivens, R. Steerenberg: "Specification of the beam position measurement in the PS machine", AB-Note-2004-001(ABP)
- [3] J.M. Belleman: "Using a Libera processor for acquiring position data for the PS pick-ups", AB-Note-2004-059 BDI
- [4] R. Garoby, S. Hancock, J.L. Vallet: "Demonstration of bunch triple splitting in the PS", CERN/PS 2000-038 (RF)
- [5] R.Garoby, S. Hancock, A. Ozturk, J-C. Perrier, J-L. Vallet: "PS machine development report: Preparation of the nominal production beam for the AD", PS/RF/ Note 99-01 (MD)