BEAM DIAGNOSTICS

Ulrich Raich
Beam Instrumentation (BI) Group
Accelerators and Beams (AB) Department CERN, Geneva Switzerland

Abstract

Most beam measurements are based on the electro-magnetic interaction of fields induced by the beam with their environment. Beam current transformers as well as beam position monitors are based on this principle. The signals induced in the sensors must be amplified and shaped before they are converted into numerical values. These values are further treated numerically in order to extract meaningful machine parameter measurements.

The lecture will first introduce the architecture of an instrument and show, where in the treatment chain digital signal analysis can be introduced. Then the usage of digital signal processing will be presented using tune measurements, orbit and trajectory measurements as well as beam loss detection and longitudinal phase space tomography as examples.

The hardware as well as the treatment algorithms and their implementation on Digital Signal Processors (DSPs) or in Field Programmable Gate Arrays (FPGAs) will be presented.

1. INTRODUCTION

1.1 Architecture of an instrument

Even though this CAS school is dedicated to *digital signal processing,* it is still a **C**ERN **A**ccelerator **S**chool and as such it should talk about techniques used in accelerators. Digital signal processing is a very technical field of electronic engineering and computer science and many people working in this field are not necessarily experts in accelerator physics. For this reason the lectures on beam diagnostics try to bridge the gap between accelerator physics and its measurements of beam parameters and the purely informatics and electronics aspects of digital signal processing.

Before looking at a few dedicated beam instruments, using digital signal processing principles, in some detail, let us first analyze the architecture of a beam measurement instrument in general. It consists of the following elements:

- The sensor (and maybe actuator) interacting with the beam
- The front-end analogue electronics
- Cabling to get the signals from the accelerator tunnel to an equipment area
- Conversion of the sensor signals to digital values
- Data acquisition and control (readout of the digital values)
- Transformation of the acquired values to humanly understandable machine parameter values
- Transfer of the results to the control room
- Graphical representation and interpretation of the results

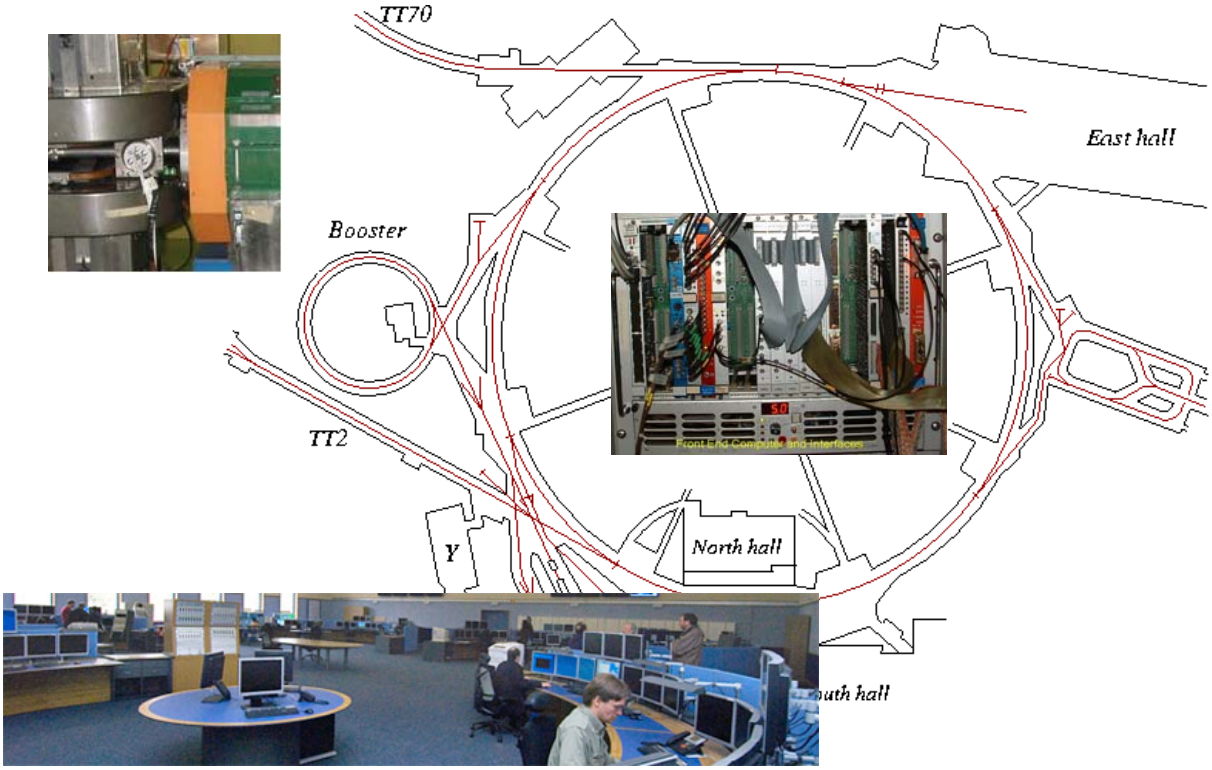

Figure 1: The architecture of a measurement system

In order to avoid signal loss and capture of electro-magnetic noise the first stage electronics must be installed as close as possible to the sensor. On the other hand it may be subject to very high radiation doses (some kGrays). Therefore the front-end electronics is installed as near to the sensor as possible but as far away as necessary to protect the electronics from radiation damage. Typical cable length from the detector to the first stage electronics is a few (1-10) meters. The signal is amplified and shaped in this first stage of the electronics. The amplified analogue signals are then transferred to the equipment gallery and digitized. In case of smaller machines like the PS, where the equipment gallery can be reached within some 100m, the analogue signals are transferred to the equipment rack and digitized there. In very big machines, like the LHC, where the distance from the underground tunnel to the equipment galleries at the surface may be of 1 km or more, the analogue signals are digitized in an equipment gallery in the tunnel und serial transmission is used to transfer the digital values to the surface. The digital results are stored in registers or memories and read out by front-end computers (called Device Stub Controllers (DSCs) at CERN). All front-end computers are connected through a local area network which also connects to operator consoles (PC-type computers with several screens attached to them) from which operators control the machine and from where they have access to measurement results in form of numbers of graphical representation.

Digital signal processing may come into play at several levels, as soon as the sensor signals have been digitized. This may be the case in the equipment gallery in the machine tunnel for big machines or in the equipment room at the surface but it may also take place only at the level of the operator consoles. Often the front-end computers are used to do the calculations, sometimes special FPGA or DSP based equipment take over this task, sometimes, for slow applications or

very compute intensive ones, the operator consoles or special number crunching machines execute the DSP algorithms.

Figure 1 shows the layout of a typical instrument in the PS. The sensor, here a beam position monitor, is connected to its analogue front-end. The amplified signal is sent to the equipment room with its VME base front-end computer ~ 150m away. There the data are read out, pre-treated and then transferred to the main control room at a distance of some 5 km.

2. TUNE MEASUREMENTS

As a first typical example, where digital signal processing techniques are used in beam diagnostics we will have a look the the machine tune Q.

2.1 What is the machine tune?

When a particle is not exactly on its design orbit it will not pass through the centre of the machine's focusing quadrupoles and therefore it will experience a focusing force, which will lead to so-called betatron oscillations. The number of oscillations around the closed orbit is called the betatron tune. This is true for horizontal as well as for vertical displacements. If the number of oscillations corresponds to an integer value, then each small imperfection in the machine will lead to a small kick, which will be repeated each time the particle passes that point. This leads to a (integer) resonance, which will result in the loss of the beam. It can easily be shown that the same is true for half integer tunes and in fact any particle for which $l \cdot Q_h + m \cdot Q_y = r$ holds where Q_h and Q_v are the horizontal and vertical tune values. The tune diagrams show these lines of instability, see Figure 2.

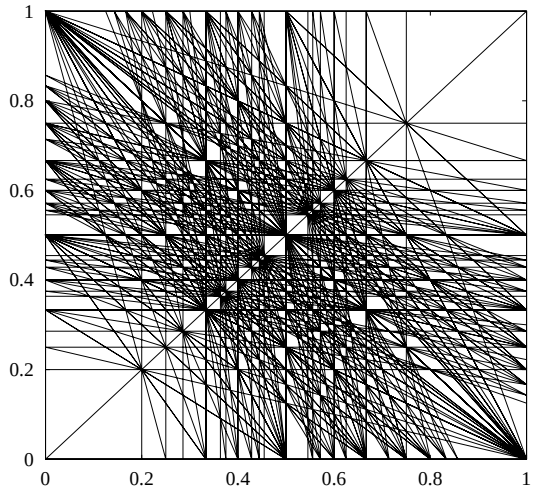

Figure 2: The tune diagram

The machine must operate in a resonance free zone of the tune diagram (the working point) in order to avoid beam blow-up and finally loss of the beam. This shows the importance for a precise determination of the tune values.

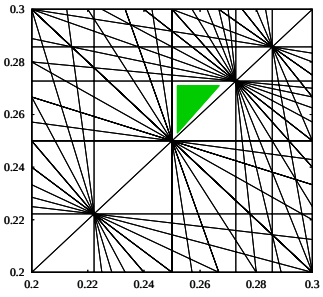

Figure 3: The working point

2.2 Measuring the tune

In order to measure the tune values, the betatron oscillations must be observed. This can be done by placing beam position monitors all around the accelerator ring and observing the beam trajectory. Since normally the beam circulates on its closed orbit with very small (unmeasurable deviations) the beam is usually coherently excited with a kicker magnet. As we have seen in the previous section we are actually only interested in the non-integer part of the tune and this can be

measured with a single beam position monitor. In Figure 4 you can clearly see that it is not possible to distinguish q from 1-q and that in addition to the base frequency higher harmonics can fitted to the data measured by the pick-up.

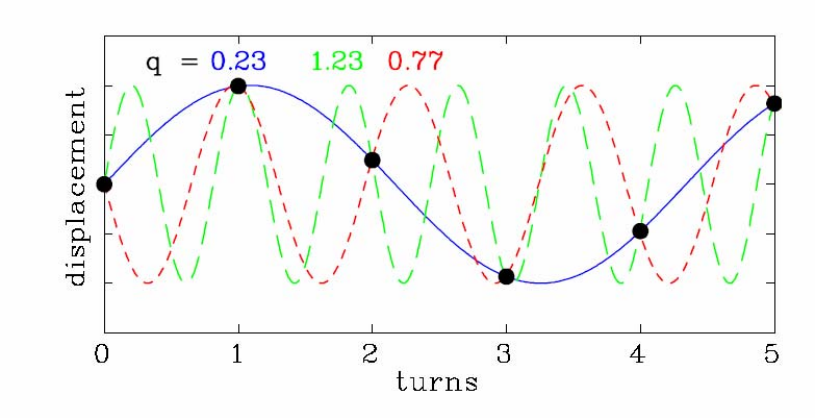

Figure 4: Measuring tune with a single pick-up

The measurement at the CERN PS and PS Booster works as follows: The beam is exited with a kicker in short intervals (min 5 ms). The beam position is measured by a pick-up (PU) delivering horizontal and vertical difference signals which are proportional to the beam displacement. The PU is of shoebox type, consisting of a metal box with a diagonal cut. In fact, the PU used here combines horizontal and vertical cuts in a single device with 4 taps to extract the signals from the left/right and upper/lower plates.
A passive hybrid circuit, being virtually insensitive to radiation, is mounted directly on the pick-up (the denominations pick-up (PU) and beam position monitor (BPM) are used interchangeably). The hybrid combines the signals from the PU plates and converts them to difference (Δ) and sum (Σ) signals.

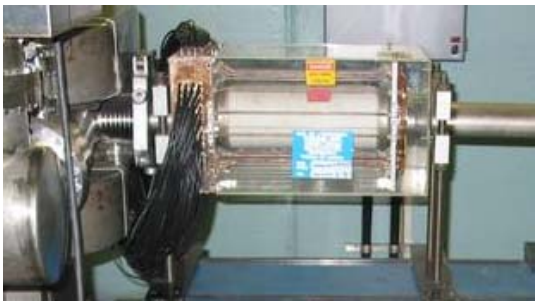

Figure 5: The tune kicker

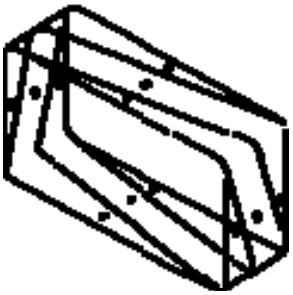

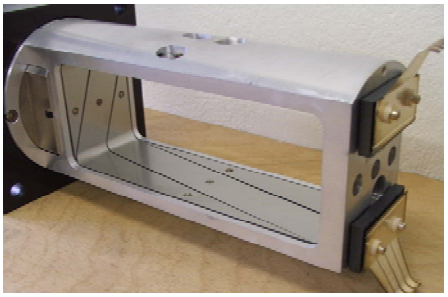

Figure 6: Pick-up electrodes

During acceleration the beam may undergo slow but rather large closed orbit variations which are filtered out by a Beam Orbit Signal Suppressor (BOSS).

Figure 7 shows the architecture of the readout electrons at the PS Booster. The Booster actually consists of 4 vertically stacked synchrotrons. The MUX is used to select the position signals from any of the 4 Booster rings. The difference signal is not only proportional to the displacement but also to the beam intensity, which may vary by several orders

of

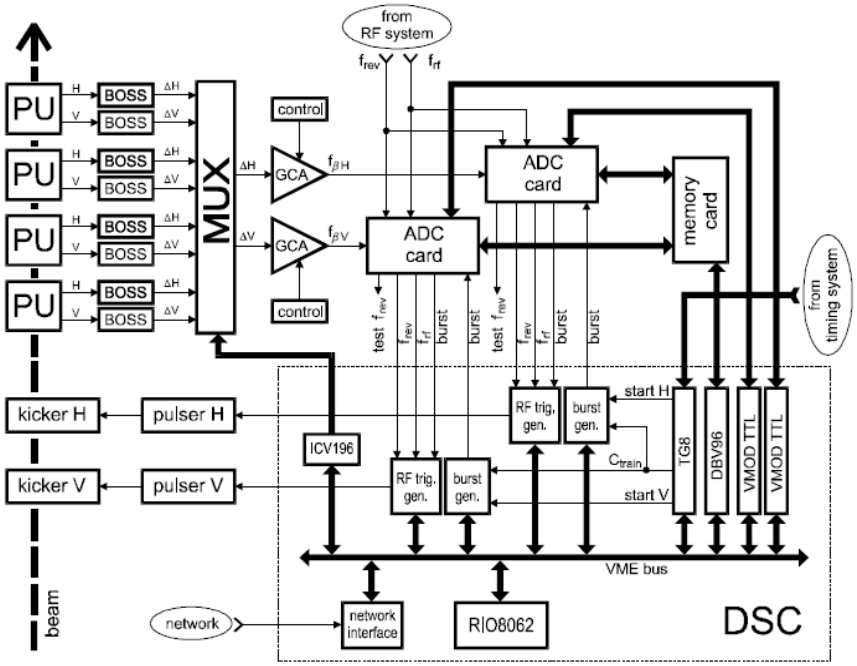

Figure 7: Tune measurement electronics

magnitude. This is taken care of by a gain controlled Amplifier (GCA). The signal is converted to digital values with a 14 bit ADC and 2048 converted values are stored in a memory card. A digital signal processor card (DBV96) with a Motorola 96002 Digital Signal Processor accesses these data over a dedicated high speed memory bus.

The results are stored in DSP memory and can be transferred to the VME processor through a communication mailbox implemented in the DSP. From there the results are transferred over the local area network to the operator consoles in the control room for visualization.

2.3 Calculating the tune

As can be seen from Figure 4 the frequency of the measured signal is related to the revolution frequency and the oscillation mode through

$$f_\beta = (m \pm q) f_{rev} \qquad (1)$$

where m is the integer number of oscillations and q is the non-integer part of the tune. Since the integer part of the tune stays constant only q needs to be measure. This reduces the equation to

$$\frac{f_\beta}{f_{rev}} = q \qquad (2)$$

As already mentioned, the closed orbit is not centered and stable during acceleration which results in an additional component of the revolution frequency in the signal. A convenient way to calculate q is to sample the PU signal at a rate proportional to the revolution frequency and to perform FFT analysis on N samples. Using a sampling rate of $k*f_{rev}$ (k is the over-sampling ratio) q can be determined by

$$q = k\frac{n_\beta}{N} \qquad (3)$$

where n_β is the bin number in the FFT spectrum.

2.4 Data handling

After starting the measurement the beam is kicked and 2048 beam positions are measured and stored in a dedicated memory card. These data are transferred to the DSP and the memory is again freed to take new data. Through this double buffering data can be taken while the previous batch is processed by the DSP.

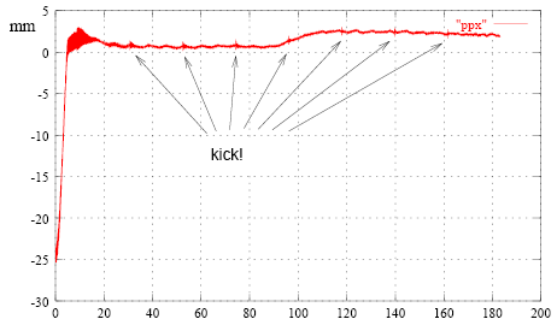

Figure 8: Kicking the beam

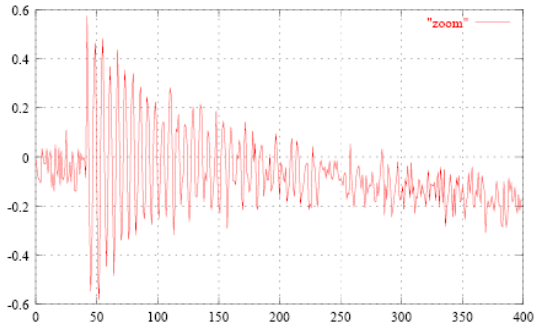

Figure 9: Zoom of the kick

The Fourier transform expects a signal that constitutes one cycle of a periodic signal. Since our signal is cut out of the stream of oscillations in an arbitrary manner, discontinuities occur when extending it to a periodic wave form.

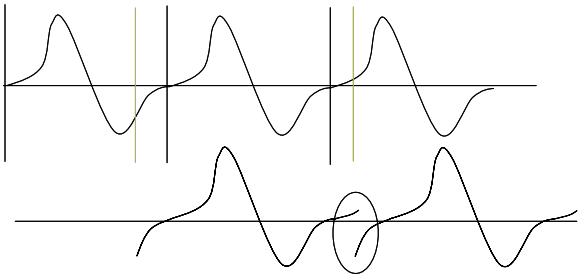

Figure 10: Periodic extension of the position signal

This phenomenon can be corrected by windowing. Each value is multiplied by a coefficient which approaches zero at the edges of the window. Typical windowing functions are shown in Figure 11. For the tune measurement system the Blackman Harris function is used. The values are pre-calculated and stored as coefficients in a table at initialization of the DSP.

For the calculation of the Fourier spectrum a 1024 point complex FFT is used, followed by an algorithm to extract the real signal spectrum from the complex spectrum data.

From the Fourier spectrum the q value must be found through a peak search routine. Since the over-sampling ratio in equation (3) amounts to 4 only a quarter of the power spectrum needs to be looked at. Q can only take only values between 0.1 and 0.5 which means that the search is limited to a window of 50 and 256.

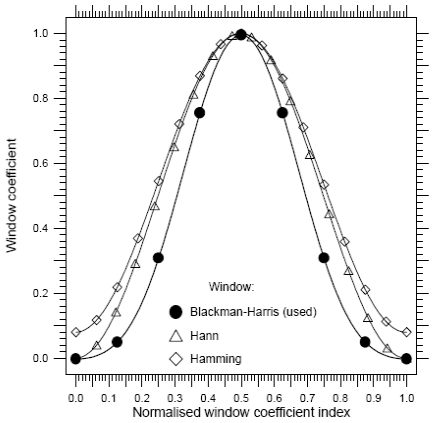

Figure 11: Windowing functions

The peak search algorithm first looks for the for bins for which

- the power value V^2 is bigger than the next value: $V^2(n-1) < V^2(n)$
- the power value V^2 is bigger than the following value: $V^2(n+1) < V^2(n)$

- the power value satisfying the first 2 conditions is the biggest in the observed range
- the power value is at least 3 times as big as the arithmetic mean of all power bins. The value 3 has been determined experimentally.

If these conditions cannot be fulfilled then the signal is either too small or the spectrum too noisy.

Most of the mathematics code for the DSP has been written in the C language except of small parts like the arithmetic mean, which has been implemented in assembler for speed reasons.

Figure 12: The frequency spectrum

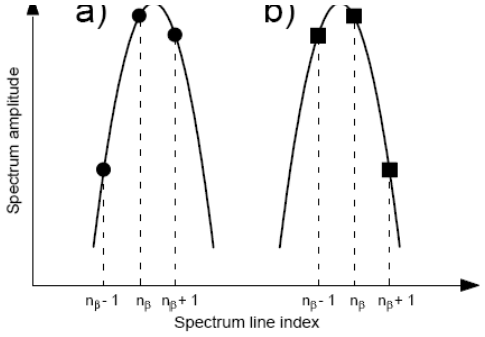

Figure 13: Tune value interpolation

2.5 Spectrum Interpolation

When just taking into account the bin in which the peak is found, then the resolution for Q is very limited. In order to improve the resolution one can

- Increase the number of points, but this would increase the computing time and therefore slow down the system
- Decrease the over-sampling, but this would lead to a degradation to the system input dynamics
- Interpolate between adjacent bins.

Since the betatron frequency spans over several bins, not only due to the discrete FFT but also because the incoming signal is not a pure harmonic but it includes attenuated revolution frequency components, noise etc. Interpolation allows finding the peak position with a resolution better than a bin. The exact form of the peak is unknown but it could be shown that using a parabolic shape, which is easy to calculate, can improve the Q-value calculations significantly.

2.6 Measurement results

During the acceleration cycle the beam is kicked every 5 ms and the Q value is calculated for each of these kicks. A front-end computer reads the results from the DSP and passes them on to an application program in the central control room, where the evolution of the tune during the cycle can be observed in graphical form. The application program also allows controlling the measurement interval, amplifier gains, start of measurement in the cycle and other parameters.

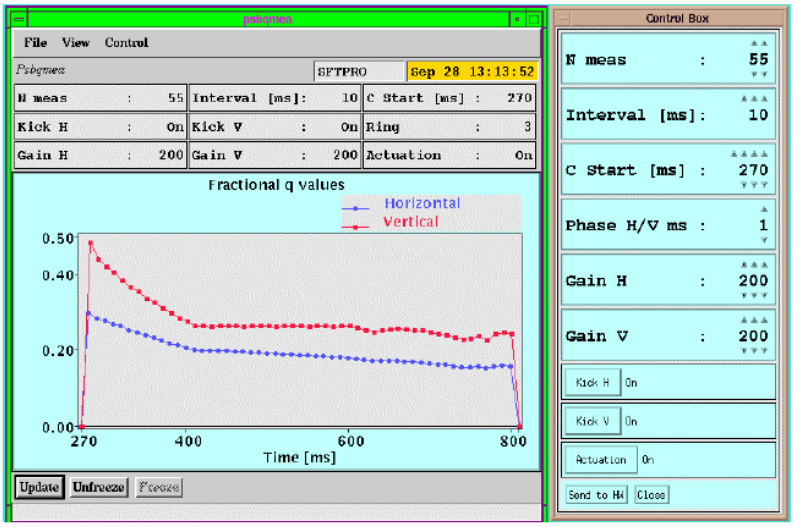

Figure 14: Tune measurement results

2.7 Further improvements

As can be seen from Figure 9 the position peaks are very sharp because the beam bunch is short with respect to the revolution frequency. As a result the signal power is distributed over many harmonics in the frequency spectrum. By rectifying the signal with a simple diode and capacitor circuit (see Figure 16) the signal spectrum can be pushed to the base band and thus the signal power in the tune line to be measured can be significantly increased. This shows that a careful balancing between analogue and digital methods is needed.

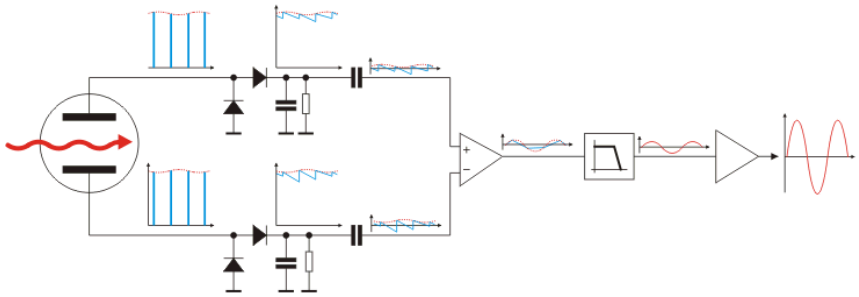

Figure 16: Base band Q measurement system

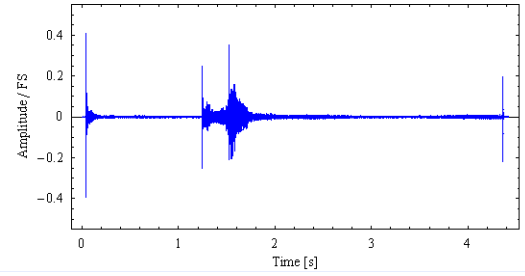

Figure 15: Natural beam oscillations

Through this trick it was possible to measure tunes in the SPS without artificial coherent excitation of the beam.

Figure 15 shows beam oscillations as they arise during an acceleration cycle in the CERN SPS. The first strong peaks are due to injection, the last peak is created when the ejection kicker is activated. Figure 17 shows a waterfall model of the corresponding frequency spectra. It can be seen that the

frequencies fall into the audia range which allows to use very powerful and cheap audio systems for data acquisition.

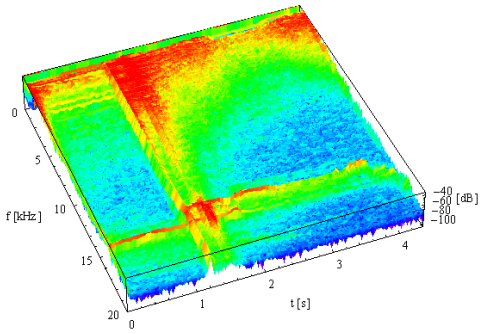

Figure 17: Waterfall model of freqency spectra

The tune measurement system described in the previous sections was state of the art some 10 years ago but it is still in operation today. The DSP chip is now obsolete and the software development system consisting of C-compiler, assembler and debugger uses PC hard-and software not available any more.

Upgrades of the digital hardware are foreseen where the DSP will be replaced by an field programmable gate array (FPGA) allowing much increased calculation speed. Like this 128000 data points can we treated instead on the 2048 and Q interpolation will not be needed any more. This shows the common problem that digital hardware and its associated software tools are evolving much faster than the accelerators which have lifetimes of several decennies (the CERN PS was inaugurated in 1955!).

3. TRAJECTORY MEASUREMENTS

3.1 Trajectories and Orbits

When particles are injected into a circular accelerator, injection errors may occur, resulting in coherent oscillations around the closed orbit. These oscillations may be caused by beam displacement and/or wrong injection angle and may cause emittance blow-up through filamentation. The oscillations can be seen when beam positions are measured turn by turn. When traversing transition (the moment when due to acceleration the orbit radius starts shrinking instead of increasing) the beam may undergo trajectory changes as well as at ejection. Other effects during acceleration may also influence the trajectories. For this reason it is desirable to be able to measure beam trajectories at any time during the acceleration cycle.

In contrast to the beam trajectory, the beam orbit is the average beam position over several turns in the machine. During acceleration the orbit may move considerably and beam losses may occur due to aperture limitations. It is therefore important to be able to measure the orbit all along the acceleration cycle.

3.2 Synchronisation

In order to measure beam trajectories it is necessary to have a big number of BPMs in the machine evenly distributed along the accelerator ring. In the CERN PS there are 40 PUs while in LHC several thousands will be installed. Since the accelerator can treat several particle-bunches (depending on the RF frequency) at the same time, each bunch must be measured turn by turn and the position data saved for the whole accelerator cycle. Another, however less flexible method, is to define the time interval during the acceleration cycle we are interested in beforehand and save the data only for this interval.

For low energy accelerators, where the particle has not yet reached the speed of light, the acceleration will result in a speed increase and therefore in an increase of the revolution frequency. This means that the integration gate, needed for integration of the BPM's Σ and Δ

signals must be continuously adapted to the revolution frequency. The current trajectory measurement system installed in the PS can measure 1 bunch during 2 turns every 5 ms and has a very complex synchronization system associated to it, following the revolution frequency in order to generate the integration gate supplied to the ADC.

In order to complicate things further, the RF harmonic number (the number of RF cycles per revolution, which is equal to the maximum number of bunches that can be accelerated) can be modified during the acceleration cycle. This allows the splitting of bunches into several bunchlets or the recombination of several bunchlets into one big bunch. In such a case, the number of integration gates must be changed on the fly.

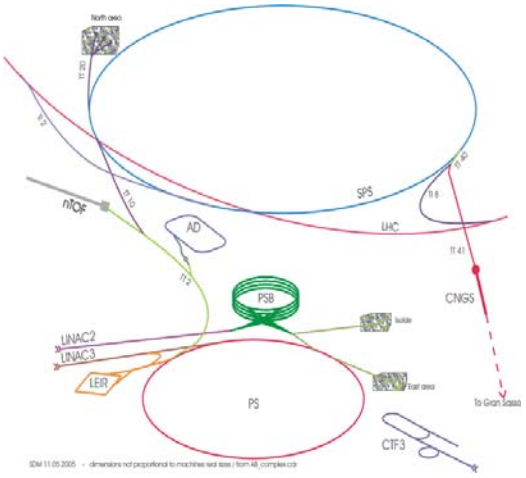

Figure 18: Different beams in the PS

The CERN control system allows several operators at different operator consoles to access data from the same measurement system. For example, one operator wants to check the orbit at injection, while another one wants to see the orbit at transition and this should be possible for the same accelerator cycle.

As can be seen from

Figure 18, the PS can take particles from Linac2 or Linac3 and it can eject particles to experimental areas (East Hall or nTOF), it can create antiprotons for the Antiptoton Decelerator (AD) or it can transfer protons to the SPS for fixed target physics or for injection into the LHC.

For maximum flexibility the PS implements a concept called *pulse to pulse modulation* (ppm) which allows assembling a series of different acceleration cycles into a so-called super-cycle which is repeated. At any moment it is possible to add or remove any of the individual acceleration cycles using a super-cycle editor. Like this several *users* can use the machine at the same time, everybody getting a few time slices of the global beam time.

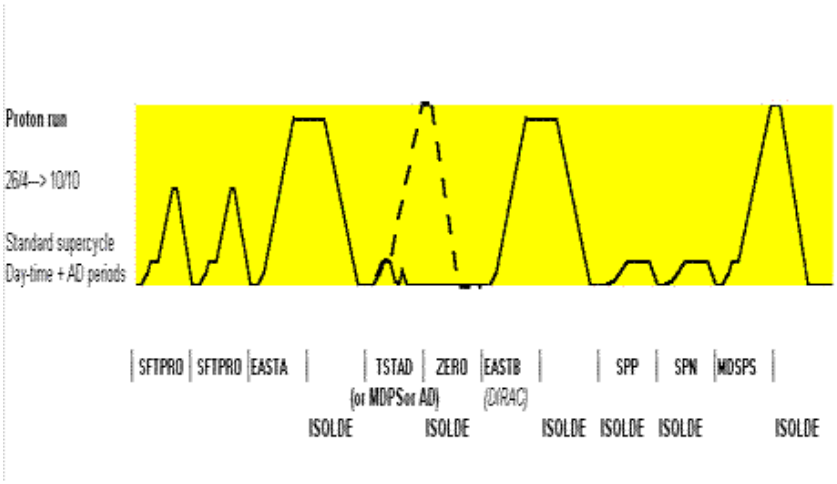

Figure 19: The PS super cycle

Figure 19 shows a typical super cycle. For low proton energy the cycle time in the PS is 1.2 s while for high energy the cycle time is doubled. For these cycles the Booster may accelerate 2 cycles where the one acclerated during the time the PS is busy is extracted from the Booster directly to ISOLDE. The first text line under the cycle diagram denotes the cycles in the PS (SFPRO, EASTA …) while the second line shows additional cycles in the PS Booster.

The consequence of this flexibility is the fact that the trajectory measurement system must deal with a great variety of different types of beams:

- EASTB: a single bunch on h=8
- AD: 4 bunches on h=8
- LHC: 4 bunches on h=7, 2 additional bunches are injected on a consecutive Booster cycle
- SFTPRO: 8 bunches fill all buckets on h=8
- EASTC: A small bunch is injected into the first RF bucket, a second, much more intense bunch is inject into bucket no. 7.
  The intense bunch is ejected at 14 GeV. The small on is further accelerated and extracted to a different area later in the cycle.

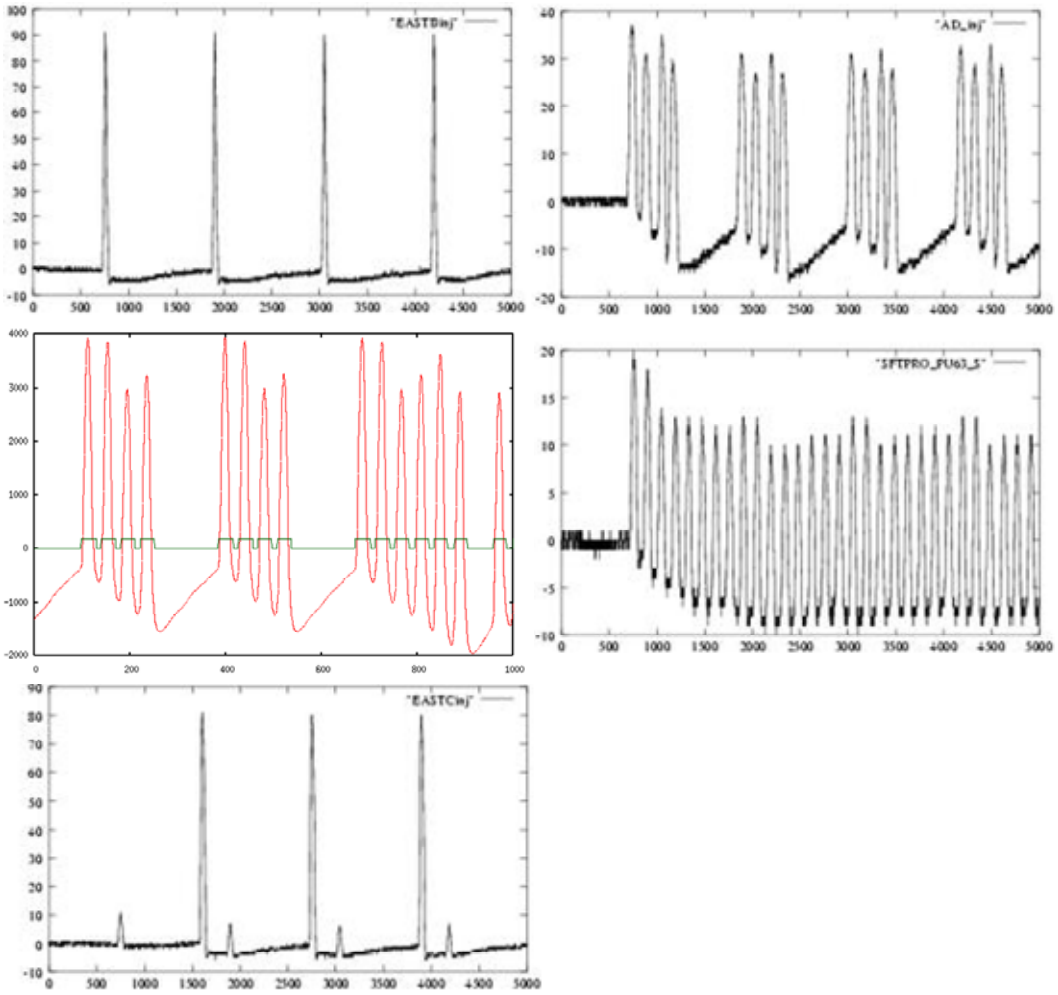

Figure 20: Beam types in the PS

3.3 Readout Requirements

In order to provide trajectory data at any time within the acceleration cycle, the bunch positions must be calculated turn by turn, from injection into the machine until ejection. This can be done by simply sampling the Σ and Δ signals coming from each BPM in the machine, fast enough to numerically integrate them and subsequently calculate the position from the integration results. It is by far easiest to use a fast sampling clock, asynchronous to the revolution frequency. With a bandwidth of the pre-amplifiers located near the BPMs of 30 MHz and expecting at least 4 samples for a 30ns bunch, we need a sampling frequency of 120 MHz. Since an acceleration cycle in the PS can take 2s, we need at least 300 MSamples per signal and therefore a total of 8 * 3 * 2 * 300Mbytes (8 bunches, 3 signals per pickup, 2 bytes per sample).

From these quick calculations it can easily be seen that it will be necessary to treat the data on the fly, such that only position data for each bunch and not individual ADC samples are stored. The idea is therefore to use an FPGA to read out the ADC samples, create integration gates synchronous to the bunch frequency on the fly and perform online baseline correction and integration. Figure 20 clearly shows how the baseline moves in dependence of the number of bunches and their intensity in the machine.

3.4 Readout Electronics

As we have seen in the previous section, the electronics must be capable of

- Reading ADC samples at a rate of ~ 120 MSamples/s
- Synchronizing to the bunch frequency
- Generating the integration gate and numerically integrate the signal
- Finding the baseline and correct for baseline movement
- Storing the integrated and baseline-corrected Σ and Λ signals in a large memory being capable of keeping all values acquired during a 2 s acceleration cycle
- Giving access to the acquired data to the external world

The only way this can be done with today's technology is the use of a fast FPGA associated with big memories.

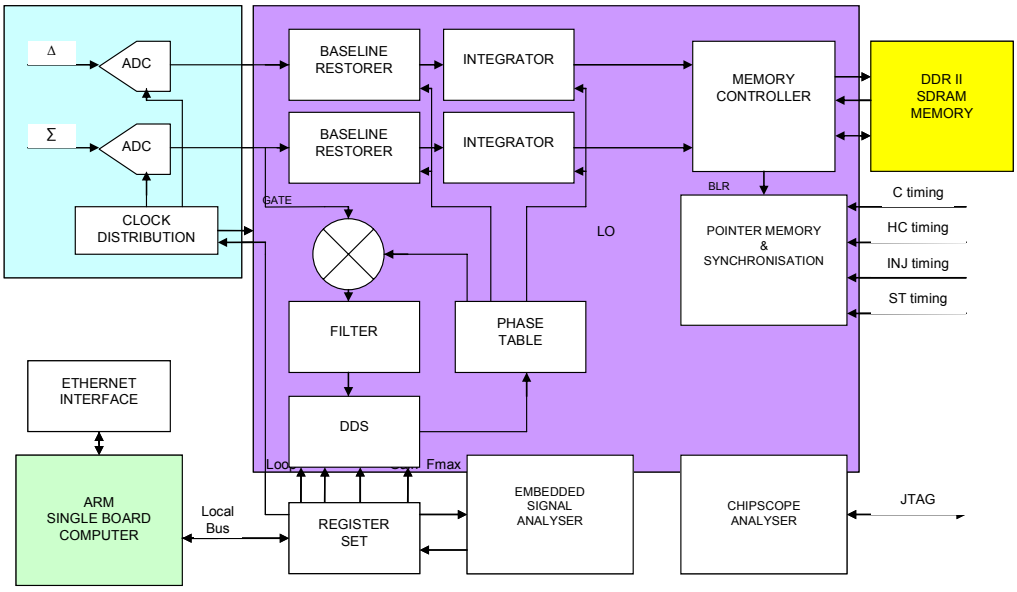

Figure 21: Trajectory measurement electronics

Figure 21 shows the basic electronics layout. The data coming from the ADC (light blue) block is treated by the FPGA (violet) and stored in memory (yellow). The data can be read out by a single board computer with Ethernet interface.

3.5 The Algorithms

In order to develop the algorithms for synchronization, baseline correction and numeric integration for the different types of beams shown in Figure 20 the FPGA is first programmed as a chart recorder, sampling passing every sample form the ADC through to the memory. Like this only parts of the acceleration cycle can be recorded but by changing the acquisition trigger any time slice in the cycle can be selected. The acquired data are read out form the memory and stored on files. The algorithms can be developed in MathLab or as offline C programs and can be tested on real data.

3.5.1 The PLL

The basic idea for synchronization is a numerical phase locked loop. The local oscillator (LO) is implemented as a direct digital synthesizer (DDS). The phase of this DDS is compared to the BPM signal and the phase error is filtered and fed back to keep the DDS synchronized. The initial guess of the LO frequency is determined from the measured magnetic field in the accelerator from which the particle energy and thus its speed and revolution frequency can be calculated.

Each incoming PU sample is compared with the phase table to decide if it is part of the signal or part of the baseline. The phase table is simply a circular buffer addressed by the highest significant bits of the DDS phase accumulator. The content of the phase table depends on the harmonic number and is pre-loaded before the start of the acceleration cycle. In case of RF gymnastics (see section 3.7) the phase table must be switched to new values on the fly.

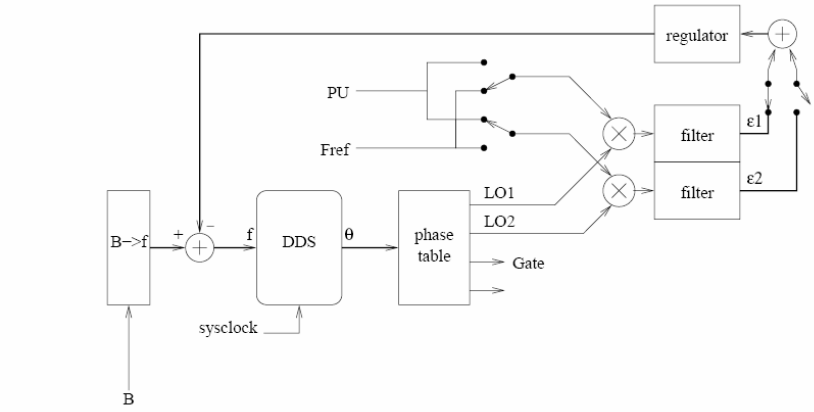

Figure 22: Phase locked loop

The problem is similar at injection: Before having beam in the machine there are no signals coming from the PUs. Instead a reference signal coming from the RF system is used for synchronization. Once the beam is circulating in the machine an external timing pulse switches to the new reference frequency.

These algorithms can be easily tested offline on the data taken with the chart recorder.

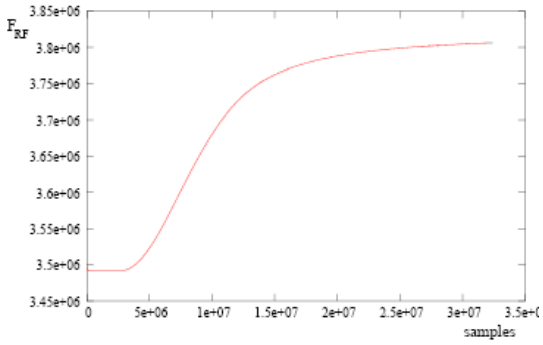

Figure 23: Measured frequency swing during acceleration

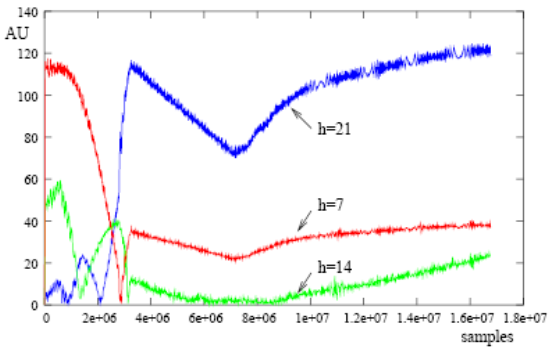

Figure 24: Frequency contents during bunch splitting

3.5.2 Baseline correction

The PU together with its load resistance delivers a high-pass filtered, and therefore differentiated version of the beam signal. The base line restorer uses a low-pass filter and therefore an integrator, with the same cut-off frequency as the PU thus compensating the differentiation effect. When integrating however the integration constant is not known and a DC level may be added to the signal. Since we know that this DC level should be zero, a second accumulator is added, which is only active when the ADC sample is part of the baseline, which can be deduced from the synchronization PLL.

3.6 Algorithm implementation in the FPGA

In the first, offline approach, the PLL in Figure 22 was implemented in a C program. The filter was a second order Butterworth low-pass filter whose coefficients were determined using a Matlab program.

$$H_F = 9.8 \cdot 10^{-6} \frac{1 + 2z^{-1} + z^{-2}}{1 - 1.9911z^{-1} + 0.9911z^{-2}}$$

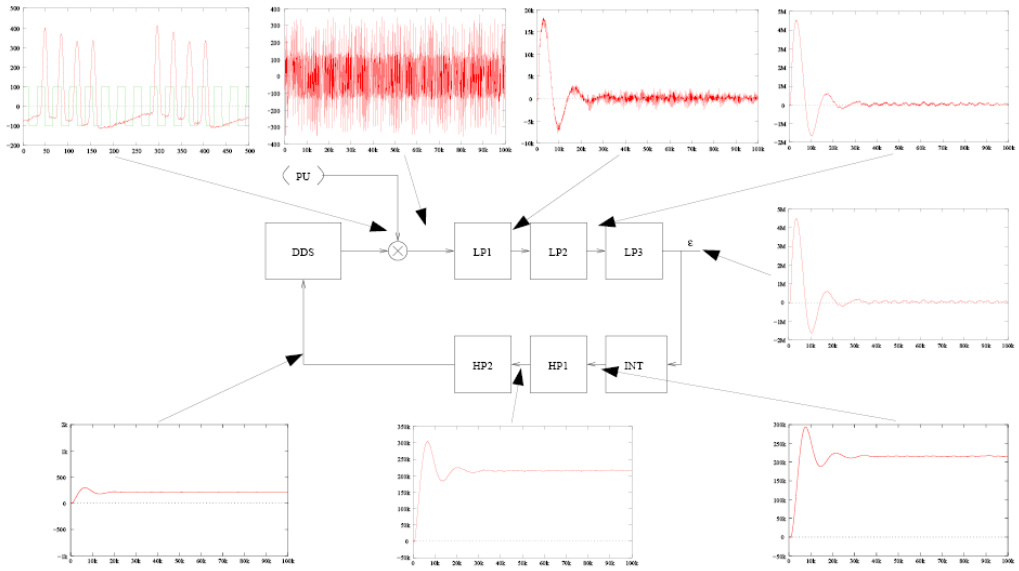

Figure 25: Splitting the loop filter

Once implemented in the FPGA the algorithms must be capable of swallowing the data coming from the PU, which means a rate of ~ 120 MHz. Calculating this formula in floating point is simply not feasible. Therefore the algorithm was re-viewed and replaced by a series of low-pass filters followed by an integrator and high pass filters using exclusively integer calculations in such a way that the coefficients can be calculated with a few shifts and additions. By subdividing the filter into several stages pipelining the algorithm in the FPGA becomes possible. Because of the large ratio between sample rate and the dominant system frequencies, a few clock cycles of pipeline delay do not affect the loop dynamics. Once the baseline is corrected and the signal is synchronized, the phase table is used to generate the integration gate. All Σ and Δ

samples within the integration gate are added and the results stored in memory. The positions can then be calculated through

$$x = S_x \frac{\Delta_x}{\Sigma}$$

3.7 Harmonic number changes

Synchronisation is largely complicated through harmonic number changes. By changing the RF frequency the bunches can be split into several sub-bunches (bunch-splitting) or the distribution of the bunches on the ring circumference may be modified (batch compression).

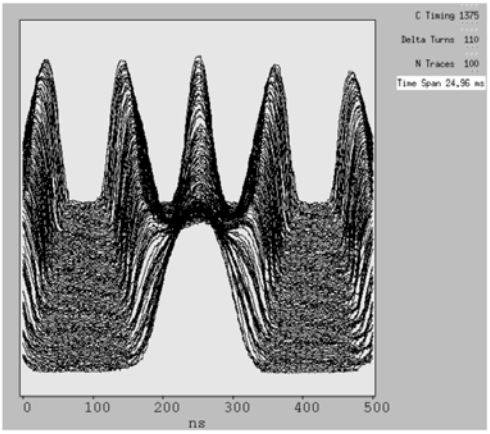

Figure 26: Bunch Splitting

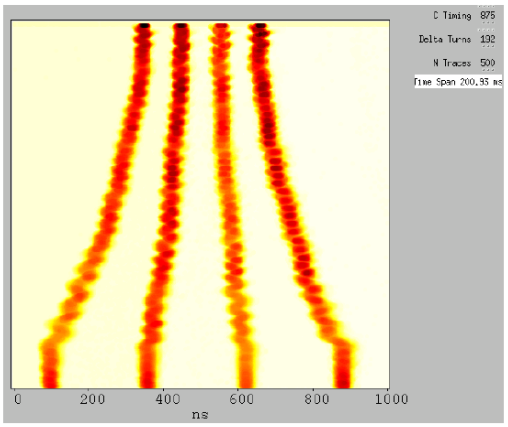

Figure 27: Batch compression

The frequency contents of the PU signals during triple splitting (a bunch is split into 3 as seen in Figure 26) is plotted in Figure 24. At a certain moment, determined by an external timing signal it must be decided to switch from measuring the position of a single bunch to measuring the positions of the 3 bunchlets. This moment corresponds to the crossing of the blue (h=21) and the red (h=7) frequency lines in Figure 24.

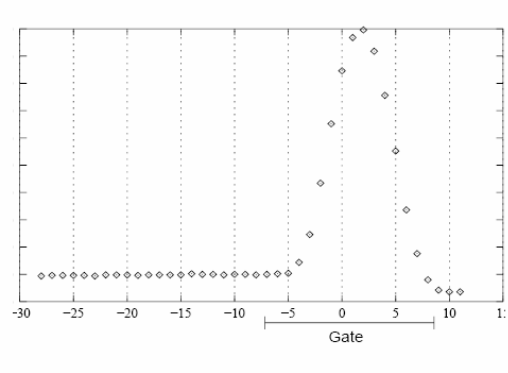

Figure 28: Integration

In order to find back the position in the table of measured Σ and Δ values where the switch has happened, address pointers, linked to these external timing events must be kept in a separate table.

3.8 Integration

Last not least the baseline corrected sum and difference signals must be integrated, which is done by simple addition of the samples within the integration gate calculated by the synchronization algorithm.

4. MACHINE PROTECTION SYSTEMS

Figure 29: Energy of a LHC bunch at 7 TeV

While in low intensity and low energy accelerators the beam cannot do much harm even if the whole beam is lost in the vacuum chamber, this is clearly not the case for high energy and high intensity machines. Already at the CERN Linac, a machine with a top enery of 50 MeV, 160 mA during a 200 μs beam pulse and a repetition rate of 1.2s it it was possible to burn a hole into one of the vacuum joints. How much more critical is the problem in the LHC, where will have 2808 bunches at 7 TeV? People amuse themselves calculating the equivalent kinetic energy of known everyday objects. The energy of single bunch in the LHC at top energy corresponds to a 5 kg bullet at 800 km/h[1] and don't forget that there are 2808 of them in the machine. On the other hand these bunches circulate in a magnetic field which is generated by supra-conducting magnets. The loss of a very tiny fraction of these particles will result in a magnet quench and the loss of a big amount, concentrated in a small volume will result in destruction of the equipment.

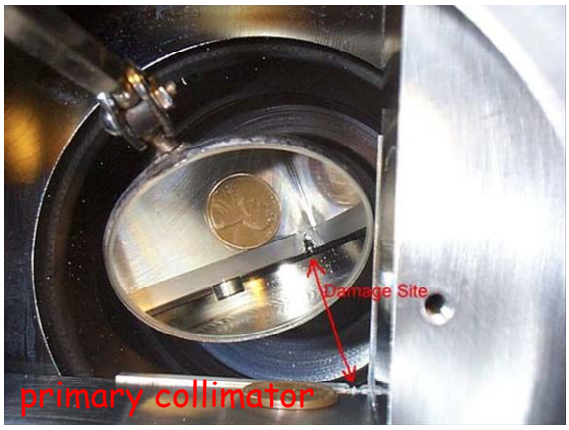

Figure 31: Beam dammage

Figure 30 Ionisation Chamber

At the Tevatron, even though its beam power is 200 times less than the power in an LHC beam, a hole was drilled into the primary collimator (Figure 31), when the beam was displaced by 3 mm due to insertion of a movable device, when it was supposed to be out of the vacuum chamber. The secondary collimator was largely damaged and 16 magnets quenched.

4.1 Machine protection using Beam Loss Monitors

The strategy for machine protection at LHC is based on the measurement of beam loss with dedicated Beam Loss Monitors (BLMs). When a high energy particle is lost, it will produce a particle shower whose energy is partially absorbed in the surrounding magnet coil but part of which can be detected by the BLMs. As long as the calibration factor of energy deposition in the magnet coils with respect to the energy deposited in the BLM is known, the BLM signals can be used to trigger beam dumps as soon as the energy deposition in the magnets gets to a level where a magnet quench must be feared. By dumping the beam and thus avoiding a quench, long beam downtime can be avoided. For even bigger beam losses, where the magnets or other equipment are at risk of destruction, this is even more true.

On the other hand the levels for triggering beam dumps must

[1] Rüdiger Schmidt, http://rudi.home.cern.ch/rudi/docs/VisitLHCWuppertal2006.ppt integration and

not be set too low because this may prevent running the machine altogether.

In order to set the trigger levels correctly the allowable losses must be known. These depend on the energy of the primary particle and on the loss duration as can be seen from Figure 32. It is therefore important to integrate the measured losses over several time periods and set corresponding thresholds for each of these integration periods. Only like this short and strong losses can be treated as well as smaller but prolonged losses.

For safety systems, in addition to the standard technical specifications like dynamic range, resolution, response time etc, the *Mean Time Between Failur*e (MBTF) is an important parameter. It defines how secure the system actually is. When the security system fails itself it must go into a failsafe state, which, as a consequence, makes the protected system unavailable.

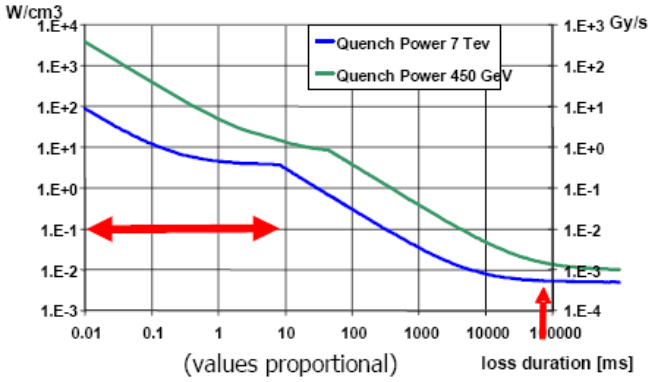

Figure 32: Quench levels

Most of today's machine protection systems use ionisation chambers as their BLMs (the LHC BLM can be seen in Figure 30). These are gas filled detectors (N_2 in case of the LHC BLMs) with the following properties:

- High dynamic range (10^8)
- Very good resistance to radiation (several MGray/year)
- High reliability and availability

The BLM signal is integrated using a charge balanced integrator and converted to a proportional frequency with a current to frequency converter (CFC). This method allows handling of a very high dynamic range. The CFC frequency is counted over a period of 40 µs. For very low currents and to allow faster response an ADC was added. The converted signal is transported from the machine, some 80m underground to the surface through optical fibres. 8 detectors are multiplexed onto an optical link and the links are doubled. The digital signal treatment is performed by a radiation resistant FPGA which:

- Reads out the converted signal
- Encodes the values in order to prepare them for transmission over the redundant optical serial links
- Multiplexes values from 8 detectors onto a signal transmission channel
- Performs CRC calculations

The electronics on the surface receives the values from the optical links, checks the CRC, de-multiplexes the signals from the 8 BLMs and gives access to the beam loss values through the VME bus. It also connects to the Beam Interlock System (BIC) generating beam dumps, should this be necessary.

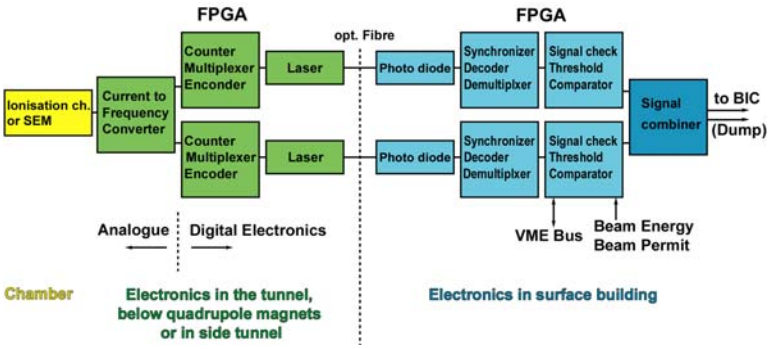

Figure 33: Layout of the BLM electronics

4.2 The Data Acquisition Board

Already several times during this lecture we have seen an electronics layout consisting of ADCs, an FPGA reading out the converted signals and treating them through fast digital signal processing algorithms implemented as VHDL code in the FPGA, followed by memory, used to store the final result. Some sort of access to this memory is needed in order to further treat the result and make it available on the operator consoles in the control room in form of easily readable numerical values, tables or graphs.

The requirement of FPGAs connected to some sort of external signal in conjunction with access to the final stored data is general enough to allow the design of a generic card which can be used for a variety of applications.

Figure 35: Photo of the Data Acquisition Board

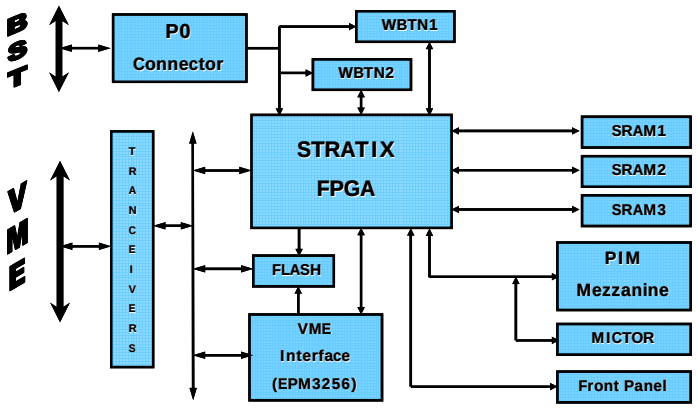

Figure 34: Block diagram of the Data Acquisition Board

The Data Acquisition Board (DAB) is implemented as a VME board built around an Altera Stratix FPGA. The FPGA can be boot-strapped through a flash memory containing the FPGA

code. This memory is also accessible via the VME bus. In addition external timing signals (Beam Synchronous Timing (BST)) can be accessed by the FPGA and used within the FPGA algorithms. In order to be as flexible as possible the FPGA has access to a local mezzanine bus. Mezzanine boards implement the application specific hardware. Like this the DAB can be used to treat signals from the orbit system, based on several thousands of BPMS, it can be used by intensity measurements systems where signals from beam current transformers are treated or it can be used to handle the communication protocol in order to read out signals from the beam loss system.

4.3 Data treatment for the Beam Loss System

4.3.1 Data treatment in the tunnel

The signal coming from the BLM is converted into frequency for coarse conversion, where a counter is used to get the final coarse value. The voltage measured on the ADC is the remainder between the last count and the first count from the next acquisition. Before sending these data to the optical communication channels, the ADC and counter data are combined to a 20 bit loss figure. Eight such values are encoded, multiplexed and sent to the surface electronics after a CRC value has been added.

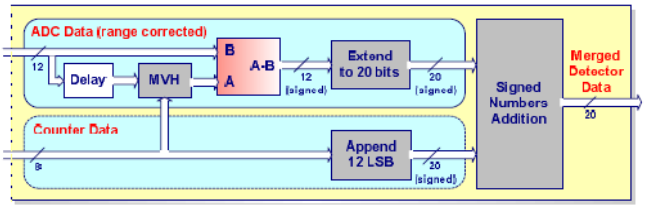

Figure 36: Beamloss dynamic range

4.3.2 Data treatment on the surface

The data treatment algorithms in the DAB, installed at the surface must perform the following actions:

- Receive the values from the electronics in the tunnel via the optical fibers
- De-multiplex the data coming from different BLMs
- Check the CRC and compare the data coming from the redundant communication channels. If the data from the two channels differ: decide which one is right
- Calculate successive sums in order to see fast big losses as well as slow small losses.
- Compare the successive sums to threshold values in order to trigger beam dumps should the losses be too high
- Give access to beam loss data for inspection in the control room together with status information
- Keep measured data in a circular buffer for post mortem analysis

- Provide error reporting

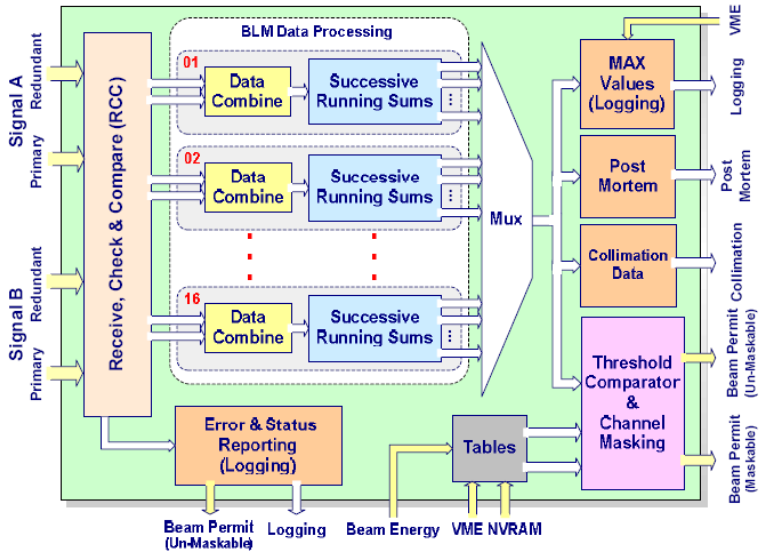

Figure 37: Algorithms treating the BLM signals

As can be seen from Figure 32 the threshold levels above which a beam dump must be triggered, depend on the beam energy and the loss duration. For this reason the loss values are integrated using running sums. 12 integration periods, spanning from 40 µs to 84 s are made available.

The calculation of running sums is rather simple. Each incoming value is added to a register, while the oldest value of the interval is subtracted again. Of course all values making up the sum must be kept.

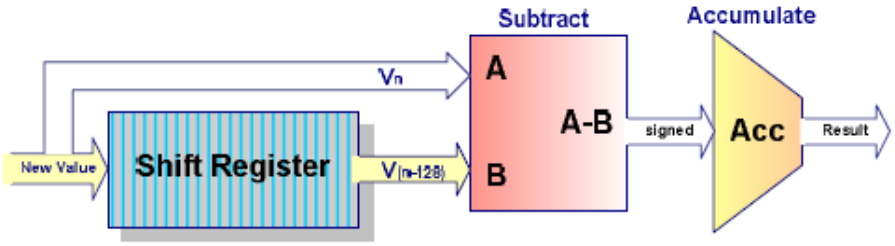

Figure 38: Calculating running sums

Another way to handle the calculations would be to use a long shift register and always add the difference between the first and the last values.

The longer the integration interval, the longer the shift register must become. This problem can be overcome by using partial sums instead of keeping all previous values. This results in a cascade of shift registers and adders as shown in Figure 39. In addition a multi-point shift register is used calculating 2 running sums at once.

Of course the latency of the sum output for each running sum depends on the time the previous sum needs for its calculation and increases therefore for longer integration periods. The running sums are compared to threshold values, dependent on the particle energy and a dump trigger is issued if the losses exceed those threshold values.

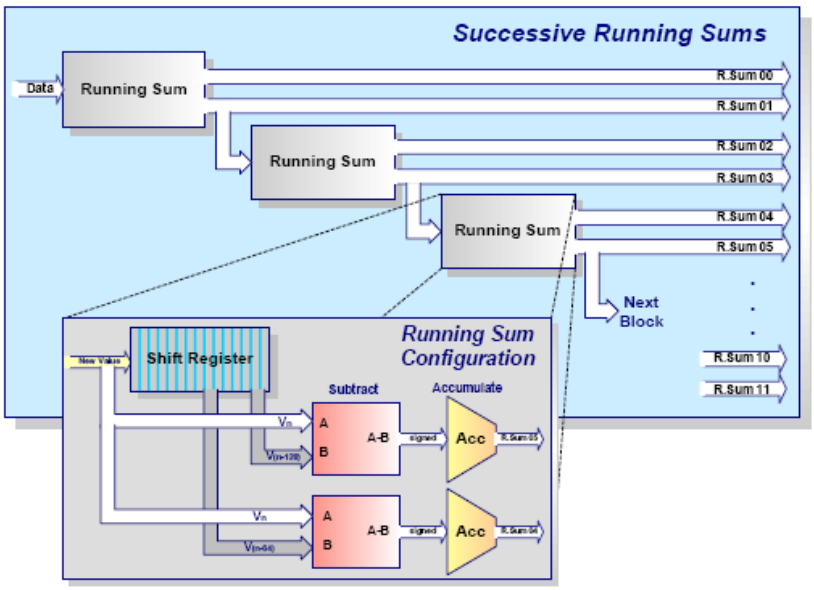

Figure 39: Cascade of running sums

Apart from triggering beam dumps the loss values are used for online viewing and logging as well as for post-mortem analysis. The values for the last 20000 turns (40 μs samples) as well as the 82 ms summed values of acquired data during the last 45 mins are available. In addition status information about the functioning of each individual BLM station is transferred from the tunnel to the surface electronics allowing online supervision of the whole system.

Part of the recorded data will be used

- to drive an on-line event display in the control room and
- write an extensive logging database both at a refresh rate of 1Hz.

Other parts of the same processing units, initiated by external triggers, will provide fast updates of the loss pattern seen. For example:

- For the automated collimator adjustments, it will record and provide the last 20ms by 640us integrals.
- At every beam injection and scheduled dump, 100ms worth of data will be pushed to the relevant systems to be used to verify the correctness of those procedures, and
- A detailed post-mortem analysis study will be possible, in the event of an unforeseen dump, by analysing the last 1.7s by 40us integrals stored in the electronics for each channel.

5. PHASE SPACE TOMOGRAPHY

Up to now we only looked at digital signal treatment where the calculations were done on the fly and which had to be very fast in order to keep pace with the incoming data flow. However sometimes it may be enough to collect the raw data and treat them using digital signal processing algorithms only after the acceleration cycle. A typical example for such a system is the measurement of longitudinal phase space distributions using tomography.

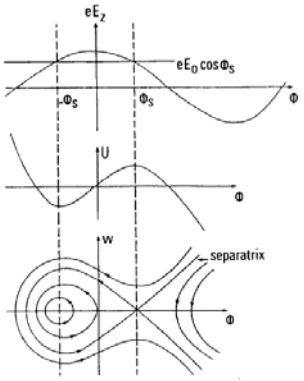

Figure 40: Longitudinal phase space

5.1 Longitudinal phase space

Consider a circular machine with a single RF cavity for acceleration, powered with a sinusoidal accelerating field. Synchronous particles having exactly the correct phase with respect to the RF field will get the pre-determined energy increase while particles coming early will be accelerated less and particles coming late will be accelerated more. This results in a movement in phase space shown in Figure 40. The longitudinal profile, measured with e.g. with a wall current monitor, corresponds to a projection of the phase space distribution onto the Φ axis.

Observing the bunch profile for many turns in the machine corresponds in fact to making projections onto the Φ axis "rotating" around the phase space distribution.

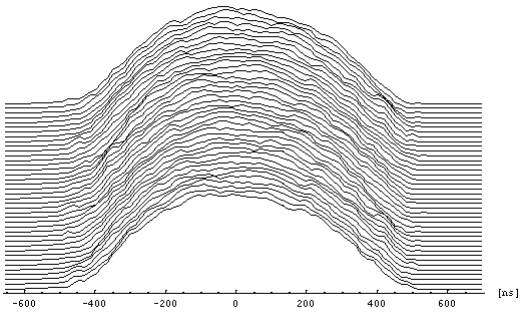

Figure 41: Bunch profiles

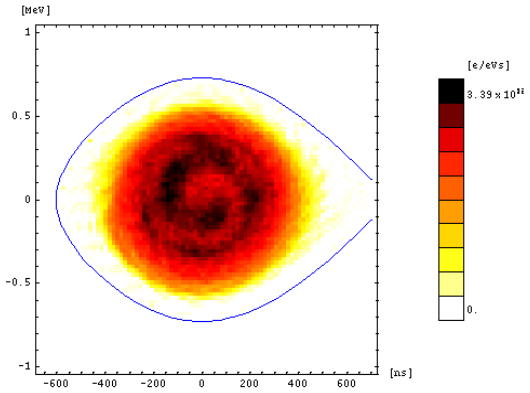

Figure 42: Reconstructed phase space

Medicine uses similar measurements to visualize 3-dimensional structures in the human body. Many X-ray images are produced with an x-ray camera rotating around the patient's body. The 3-dimension object is re-constructed using Algebraic Reconstruction Techniques (ART).

The principle is rather simple (the implementation however can be rather complex and requiring a lot of computing power): The x-ray image, which is in fact a projection of the 3-dimensional object onto the camera axis, is back-projected. This means that the content if a one-

dimensional bin is distributed over a 2-dimensional array of cells that could have contributed to that bin. A first reconstructed version of the original object is obtained. Projecting this reconstructed object and comparing the result to the projection of the original object yields differences. These differences can again be back projected yielding an improved reconstruction. Doing this for many images allows the re-construction or the original object.

Figure 43: CT scanner

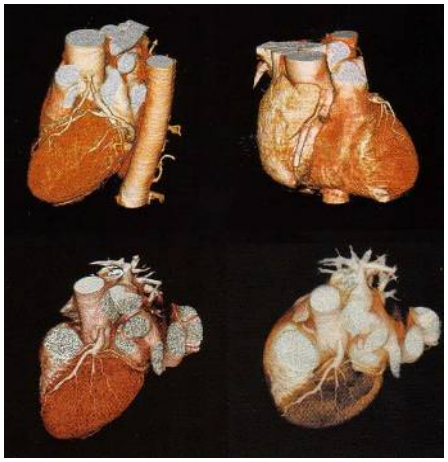

Figure 44: Results from CT reconstruction

The method of ART can be applied to the profiles obtained from wall current monitors measuring bunch profiles over many turns. While the individual measurements do not exhibit any detailed phase space information, the reconstructed phase space plot display a wealth of interesting details which cannot be observed otherwise.

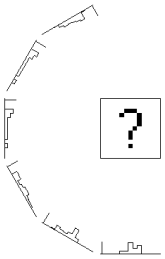

Figure 45: Projection

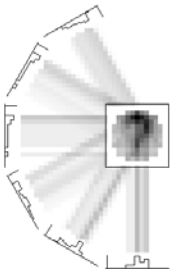

Figure 46: Back projection

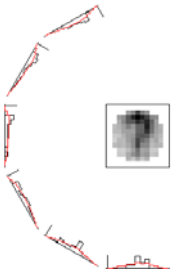

Figure 47: Projection of back projected image

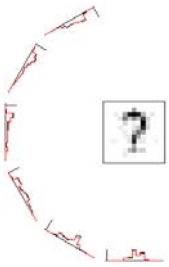

Figure 48: After 50 interations

5.2 Particle tracking

The reconstruction technique explained in the previous section only works as long as the projections are taken with a constant rotation of the viewpoint: the object is rotated in front of the camera at always the same rotation angle, which would correspond to a circular movement in longitudinal phase space. Unfortunately however the synchrotron motion of particles represents

large non-linearities. These non-linearities are taken into account by tracking particles in a simulated machine in order to build maps describing the evolution in phase space. These maps are taken into account for the reconstruction.

5.3 The Hardware

When a beam travels through the vacuum pipe an image current, equal in intensity with the primary beam but traveling in opposite direction, is created. By creating a gap in the vacuum chamber and bridging it with resistors, this image current can be measured. Such a wall current monitor (WCM) can have a very large bandwidth (typically in the GHz range) and the longitudinal shape of the particle bunches can be measured. The WCM signal is amplified and measured with a fast oscilloscope.

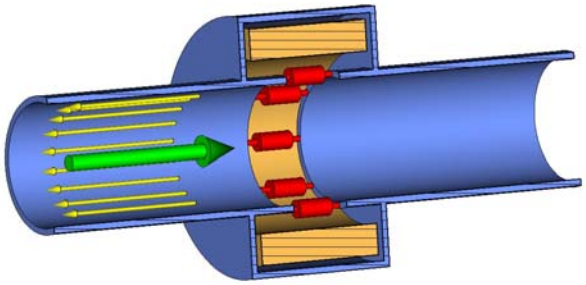

Figure 49: The Wall Current Monitor (WCM)

Of course the signal must be measured on a turn by turn and bunch by bunch basis which requires an oscilloscope trigger synchronous to the revolution frequency. The problem is similar to the one described in the chapter on trajectory measurements.

A front end computer is dedicated to the trigger timing while the scope trace data are directly read out by the application program via a dedicated (GPIB) Ethernet link thus improving data transfer speed.

The bunch profiles (see Figure 41) are used as input to the ART algorithm. In a first version the ART code was implemented in Mathematica™ because of its large mathematical libraries. Later however it was implemented in High Performance Fortran-90 which allow parallelization and execution on several processors, further reducing execution time. Many tricks were played like the use of integer arithmetic instead of floating point when this was possible as well as extensive use of tables, e.g. for the calculation of trigonometric functions. Like this the execution time could be reduced from hours to some 15s-20s.

CONCLUSIONS

Beam diagnostics is a field spanning a great number of disciplines: Machine physics, analogue and digital electronics and computer science. Quite naturally digital signal processing methods are employed to extract physically meaningful and easily understandable results from the raw signals created with the beam instrumentation sensors. Using a few typical measurement applications as examples these lectures try to bridge the gap between the methods employed in digital electronics and the beam physics profiting from the power of these methods.

6. REFERENCES:

[1] M. Gasior, J. Gonzalez, DSP Software of the Tune measurement System for the Proton Synchrotron Booster, CERN PS-BD Note 99-11

[2] M. Gasior, J. Gonzalez, New Hardware of the Tune Measurements System for the Proton Synchrotron Booster Accelerator

[3] J. Belleman, Using a Libera Signal Processor for acquiring position data for the PS Orbit Pick-ups. CERN AB-Note-2004-059

[4] J. Belleman, A New Trajectory Measurement System for the CERN Proton Synchrotron, Proceedings of DIPAC 2005 Lyon

[4] B. Dehning, Beam Loss Monitor System for Machine Protection, Proceedings of DIPAC 2005 Lyon

[5] C. Zamantzas et al, The LHC Beam Loss Monitoring System's Surface Building Installation CERN-AB-2007-009 BI, presented at LECC 2006 – 25-29 Sep 2006 – Valencia/SP

[6] C. Zamantzas et al, An FPGA based Implementation for Real-Time Processing of the LHC Beam Loss Monitor System's Data, presented at IEEE NSS 2006 – Oct. 29 / Nov. 4 2006 – San Diego/USA