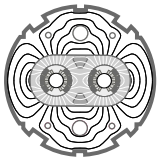

LHC Project Note 278

2002-01-22

Thijs.Wijnands@cern.ch

Prototyping LHC orbit control

Author(s) / Div-Group: Thijs Wijnands, Joerg Wenniger, Bala Srinivasan[†]

Keywords: Orbit, SVD, time constants, PID feedback control

Summary

Orbit correction consists in adjusting the strengths of the corrector magnets to make the measured beam position match a predefined reference. In the LHC, this involves around 2000 sensors and more than 1000 actuators that are distributed along both rings. The orbit correction scheme should be able to compensate for very slow orbit drifts in the range of a 10^{-2} Hz but also for fast motions (vibrations) up to 1 Hz. In this paper we investigate correction schemes that could be used in either case. The choice of design formalisms is based on the experience we gained with the SPS and the LEP.

1. Introduction

Orbit control will play a key role in the LHC due to the tight aperture. Depending on the emittance of the beam injected from the SPS, one can expect bad life times and loss of beam current when the orbit is not properly adjusted. Eventually, this will lead to a reduction of the integrated luminosity and machine efficiency.

Until experimental data becomes available, the exact frequency spectrum of orbit distortions in the LHC will remain unknown. In this paper we will therefore base ourselves on the experience gained in LEP [1]. Over the duration of a LEP fill, the orbit was kept as close as possible (tolerances on the orbit RMS were 50-100 μm) to the reference by a slow integral feedback control loop. The loop applied corrections in both planes at a maximum rate of about 2 corrections/minute. While the *horizontal* orbit was very stable and the drift almost never exceeded 0.2 mm over the duration of a complete fill, the vertical movements were fast, with integrated orbit drifts exceeded 4 mm RMS. The major part of the vertical orbit drifts in LEP was due to the movements of the superconducting low-β quadrupoles. The effects of those quadrupoles were enhanced by their large strength and betatron functions. Most of the vertical orbit correction was done using the orbit correctors installed next to the low-β quadrupoles. In the frequency range of 1 to 100 Hz, the orbit oscillation amplitudes were limited to 10 microns or less. Similar figures have been reported by HERA[2].

We expect the spectrum for the orbit motions in the LHC to be somewhat different from the one observed in LEP. There will be slow drifts (10^{-2} Hz) due to the ground motion [2] (as in LEP) and due to the decay of persistent currents at injection [3]. Movements at approximately

[†] Present affiliation : Institut d'Automatique, Ecole Polytechnique Fédéral de Lausanne, CH-1015, Switzerland

0.1 Hz are predicted at the start of the ramp due to the "snap-back" [4] and possibly during physics due to the beam-beam effects [5].

To compensate slow orbit drifts at frequencies up to 1 Hz, we envisage a central process that will monitor the orbit continuously and applies corrections whenever the root mean square of the orbit drift exceeds a certain threshold. In this paper, we investigate the orbit correction schemes that could be used in such a global feedback loop. The same principles also apply to local orbit correction. We will concentrate on the correction algorithm, the data communication and the feedback control algorithm. The approach is general and could be envisaged for other feedback applications such as the tune loop, signals from the reference magnets or chromaticity control.

2. Correction algorithm

Suppose that the vector $\boldsymbol{X}$ (size N) represents the beam position measured at the BPMs and the vector $\boldsymbol{Y}$ (size M) the corrector strengths. The task of orbit correction is to find a set of corrector strengths that will minimise the difference between $\boldsymbol{X}$ and a "golden" reference orbit named $\boldsymbol{X}^{gold}$:

$$(\boldsymbol{X}-\boldsymbol{X}^{gold}) + \mathrm{A}\boldsymbol{Y} = 0 \qquad (2.1)$$

The response matrix A (dimension NxM) describes the relation between the corrector kicks and the beam position changes at the BPMs. In general, A is not a square matrix and cannot be inverted. For the most common situation where $N>M$, the system is over constrained and a correction algorithm will search for a least square solution by minimising $\|(\boldsymbol{X}-\boldsymbol{X}^{gold})+\mathrm{A}\boldsymbol{Y}\|^2$. The problem can be more complicated if the matrix is singular or (mathematically speaking) close to singular. This can be caused by design or when some monitors are not available.

The two most popular correction algorithms are MICADO [6] and Singular Value Decomposition (SVD) [7]. MICADO is an algorithm based on Householder transformations that performs an iterative search for the most effective corrector. In the context this note, we will consider only SVD because it can be cast in a very simple form. Yet it should be clear that it is straightforward to replace SVD by MICADO. The SVD algorithm allows us to write matrix A as the product of three matrices :

$$\mathrm{A} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^T \qquad (2.2)$$

$\boldsymbol{U}$ and $\boldsymbol{V}$ are orthogonal and are obtained from an eigenvalue analysis of the covariance matrix of A. Matrix $\boldsymbol{S}$ is a diagonal matrix containing the singular values s_i :

$$\boldsymbol{S} = \begin{pmatrix} s_1 & & & \\ & s_2 & 0 & \\ & & . & \\ & 0 & & . \\ & & & s_n \end{pmatrix} \qquad (2.3)$$

The pseudo-inverse matrix Â of A is given by :

$$\hat{\mathrm{A}} = \boldsymbol{V}\mathrm{diag}(\boldsymbol{1/s_1,\ldots})\boldsymbol{U}^T \qquad (2.4)$$

where Â can be made non-singular by zeroing all elements of $\boldsymbol{S}^{-1}$ with $s_i < \varepsilon$, a predefined cut off. The solution of the least square problem is now :

$$\boldsymbol{Y} = - \hat{A}\boldsymbol{X} \qquad (2.5)$$

A priori it is sufficient to calculate Â only once before beam is injected. In practice Â must be updated regularly to keep track of faulty monitors and orbit correctors. When there is circulating beam, the task of the feedback loop consists of performing (2.5) in real time.

The eigenvalues s_i cover 2 to 5 orders of magnitude for large machines such as SPS, LEP or LHC (figure 1). The weights s_i are approximately proportional to the RMS of the orbit response corresponding to its associated eigenvector (the columns of matrix $\boldsymbol{V}$). Correction of large-scale orbit patterns is done with the large eigenvalues while smaller values correspond to smaller structures such as local bumps. The number of eigenvalues retained will therefore determine the spatial resolution of the correction. Correction of orbit drifts with respect to the "golden" reference orbit can in general be done with a few (around 20) large eigenvalues. When the number of eigenvalues is increased, the orbit deviation with respect to its reference can be reduced but it also makes the orbit correction more sensitive to noise on the BPM readings. A frequently used solution is to use only a restricted set of correctors and monitors, chosen such as to sample the overall orbit features well enough. In this way the size of the linear system can be reduced significantly.

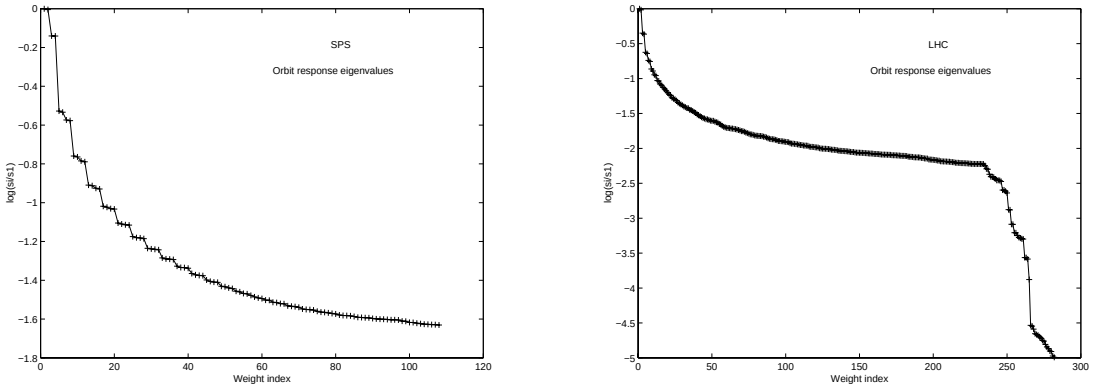

Figure 1 : SVD eigenvalues of the SPS (left) and LHC (right) vertical orbit. The very small s_i with $s_i/s_1 < 10^{-4}$ correspond to almost singular solutions in the low-β regions of the interaction regions. No such singularities appear in the very regular lattice of the SPS.

The vector of correction kicks $\boldsymbol{Y}$ has size N (approximately equal to 250 for a single beam pipe and a single plane). It is possible to apply the control algorithm to this kick vector or to the eigenvectors that represent an orthogonal basis.

3. Architecture

Similar as in LEP, a centralised solution for orbit correction is envisaged where a single processor receives and sends data to equipment front ends that are distributed around the ring. The actuators of the system are the orbit corrector magnets and the associated power converters (approximately 1000 in total). The power converter controllers are accessed via a deterministic fieldbus. The sensors are BPMs that measure the transverse position of a longitudinal slice of the bunch train. A single BPM has 4 button monitors (2 per plane) and 8 BPMs will be connected to a single BPM processor front end (approximately 1000 in total). The present estimate is that each BPM will produce 5 bytes of data at a rate of 10 Hz. For 500 BPMs per ring, this corresponds to a datastream of 50 kBytes/s per beam pipe when data is produced at a rate of 10 Hz.

4. Dynamics of the system

4.1 Time delay and time constants

Time constants and time delays are important in a feedback loop because they limit the performance of the controller and determine the robustness of the system. In our case, the plant, (i.e. power converter and associated magnet, the beampipe and the overall time delay) induce the entire dynamics we observe.

The power converter and the associated corrector magnet have a natural time constant of L/R where L is the self-inductance of the superconducting corrector magnet and R the warm cable resistance. For the superconducting corrector magnets, the L/R time is rather long (about 200 s, open loop bandwidth of about a mHz). However, during the injection, the bandwidth of the system can be increased because there is more power available than required for steady state operation.

The beam pipe acts as a low pass filter and thus introduces another time constant. However, it will be shown in section 5 that it plays only a minor role in the design of the loop.

Finally there is the total time delay between sensors and actuators. It is the sum of the time needed for data acquisition, data communication across the site and computing. A delay in a feedback system leads to a phase shift and this will *always* degrade the stability and the damping of the system. Within limits, time delays can be compensated for by adjusting the feedback gain and/or by using lead or lag compensation. For orbit control, the total time delay is rather large and it limits the performance of the feedback system.

4.2 Time constant of power converter and corrector magnet

The cold LHC orbit corrector magnets MCBH(V)[†] have an inductance of 7 H and a warm cable resistance of 30 mΩ which yields L/R = 230 s as the natural system time constant. A RST type control algorithm is driving the 4-quadrant power converter (V_{max} , I_{max} = ± 8 V, ± 60 A). This system is using the available power to accelerate the response of the system [8].

The steady state system load is at most V_{stat} = 60 A x 30 mΩ = 1.8 V when the corrector is set to its maximum current, but at injection this value will never be reached (otherwise it would not be possible to ramp). For large kicks (20 μrad or 1 A at 450 GeV) the response of the system power converter & magnet can be accelerated by a factor V_{max}/V_{stat} = 4 which gives a time constant of 58 s. For small kicks, (2 μrad or 0.1 A at 450 GeV), the ratio is V_{max}/V_{stat} = 2000 which gives a time constant of 0.1 s.

The maximum kick strengths during injection and in coast will determine the effective time constant of power converter and magnet. The most important variations of the orbit are expected to occur during the snap back at the start of the ramp and during the squeeze. In the H-plane and for a random field error b_1 of 0.75 units (see [9]), the orbit RMS change is 800 microns (covering a range extending from 250 to 1750 microns). Using an SVD correction with 50 eigenvalues, the RMS can be reduced to 70 microns (figure 2). The associated kicks are ranging between 0 and 2 μrads with a kick RMS of 0.5 μrad and evolve during the snapback.

During acceleration and coast, there is less additional power available. Around 0.2 Volts are needed during the ramp (inductive term). Up to ~ 2 Volts are needed to compensate the

[†] Note on the warm orbit correctors : in IR3 and IR7 there are 2 warm correctors per plane and per beam, which is insufficient for a local orbit correction. Including cold correctors means that even local orbit correction is dominated by the time constant of the cold corrector magnets.

Ohmic dissipation in the warm cables in coast (steady state load due to the resistive term). At top energy, the effect of the same kick is reduced by a factor 20 compared to injection energy. To get the same beam deflection the kick strengths thus have to be about 20 times stronger. As a consequence, the effective time constant of the power converter & corrector magnet is now about 60 s at the end of the ramp and during physics. Consequently the feedback system has no gain for errors at 1 Hz. However, it is still possible to use the loop provided a gain-scheduled controller is used. In this scheme, the beam energy is used as an input and the feedback gain is reduced as the energy is increased.

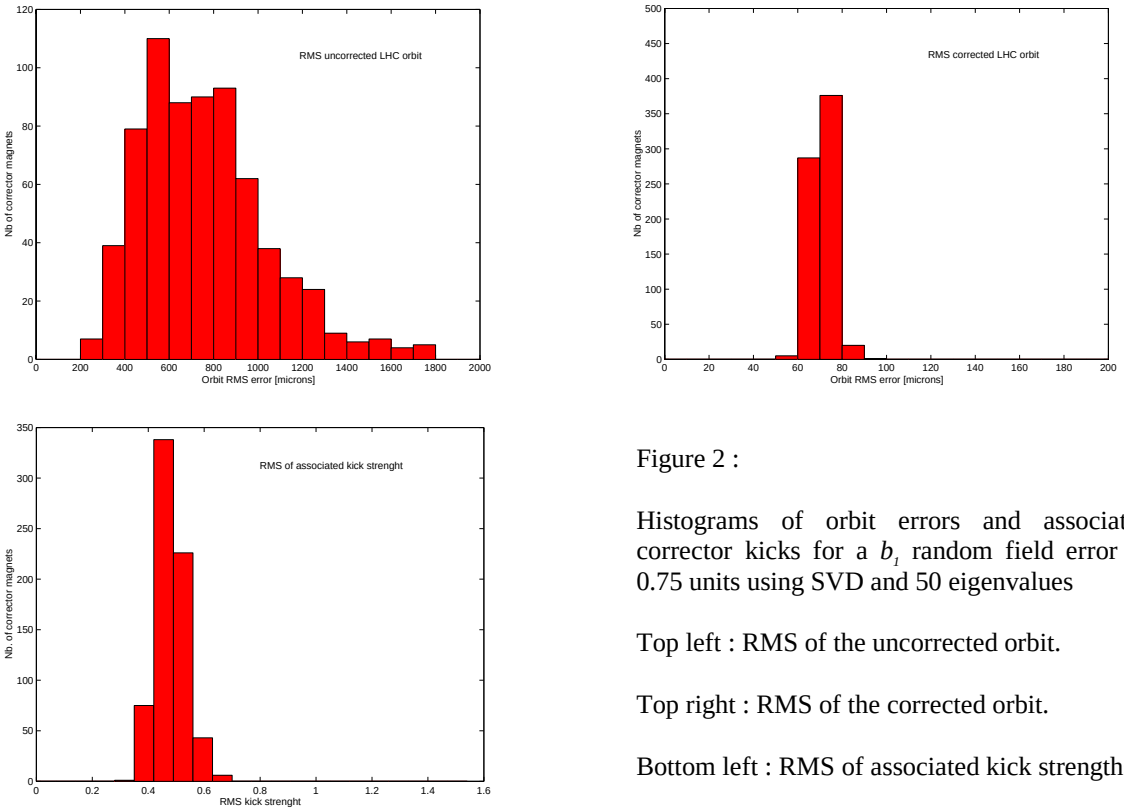

Figure 2 :

Histograms of orbit errors and associated corrector kicks for a b_1 random field error of 0.75 units using SVD and 50 eigenvalues

Top left : RMS of the uncorrected orbit.

Top right : RMS of the corrected orbit.

Bottom left : RMS of associated kick strengths.

4.3 Time constant of beam pipe

The response of any magnet will be affected by eddy currents in the vacuum chamber, which act as a low pass filter with a first order transfer function. The characteristic time constant is given by [10] :

$$\tau_{ec} = \frac{\mu_0 \pi b d}{2\rho} \qquad (4.1)$$

where b is the radius of the vacuum chamber (2 cm in the cold regions, 4 cm in the warm regions), d the thickness of the copper plating (50 μm cold and 0.85 mm warm) and ρ the resistivity ($1.7\ 10^{-10}$, RRR =100 cold, $1.7\ 10^{-8}$ warm). This yields a time constant of 12 ms for cold and 2 ms for warm chambers. Note that the time constant of the beam pipe is negligible compared to the time constant of the power converter and the corrector magnet (around 100 ms).

4.4 Total time delay

There will be a time delay between actuators and sensors due to the transport of data across the LHC site via the computer network. A first guess of the delay can be made using measurements from the SPS. A thousand 1 kByte data packets were send from building BA2 to

building BA5 using identical end nodes (Power PCs at 200 MHz running LynxOS operating system). Figure 3 (left) shows the time delay when using TCP/IP. TCP is a protocol that was designed to provide a reliable flow of data between hosts. It is using timeouts, checksums and acknowledgements to throttle the data flow and to ensure that data is received on the other end. Reliability is provided in a non-predictive manner and it is difficult to establish a data flow with a fixed time delay. Most data packets are delayed by 2.3 ms but data can occasionally arrive much later.

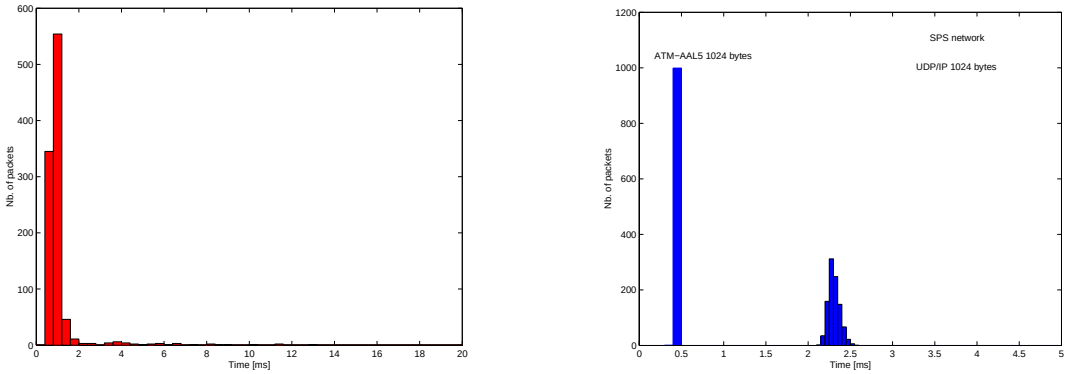

Figure 3 : Histograms of time delays when sending 1000 one kByte packets from BA2 to BA5 via the SPS machine network on 200 MHz Power PCs with LynxOS. Left : using the TCP/IP protocol. Right : using UDP/IP. The figure on the right also shows the time delay for ATM-AAL5 [11].

Figure 3 (right) shows the time delay for UDP/IP, which is a much simpler protocol than TCP. It just sends the data from one host to another without verifying whether the data has arrived. The jitter on the time delay is thus much smaller (2.3 ms time delay with a jitter of 10 % in figure 3).

Another time delay is due to execution of the SVD (or any other) correction algorithm. The algorithm retained 20 singular values out of 500 for corresponding to a single plane and one beam pipe. Figure 4 shows histograms when executing algorithm 100 times on different types of CPUs. The years in the figures indicate the time when these CPU's were first used at CERN.

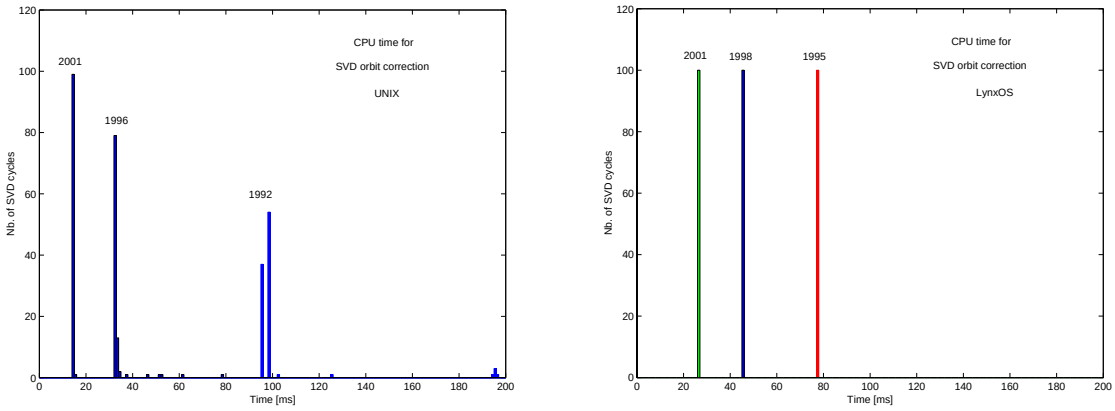

Figure 4 : Histograms of execution times for 100 SVD orbit corrections (500 BPMs and 300 correctors) on a various types of HP-UX Unix workstations (left) and on various types of LynxOS Power PCs (right). The years in the figures indicate the time when the CPU's were first used at CERN. Note the spread in execution times on UNIX and the absence of this when using a LynxOS.

The execution time depends on many board specific issues such as the processor speed, the compiler and the operating system. Note also that on a time-sharing operating system such as Unix, the CPU time varies as a function of the machine load (figure 4, left). A real time system such as LynxOS allows reservation of system resources and does not show such

variations (figure 4, right). For one of the most performing Power PCs that is presently used in operations, the execution time of a single SVD cycle for one beam pipe and a single plane is about 20 ms.

Finally, there is a time delay induced by the power converter control system. The new corrector reference is communication via a gateway and a fieldbus to the power converter controller. The estimated minimum induced delay is 20 ms-50 ms.

Delay source	***Min* [ms]**	***Max* [ms]**
Data Acquisition	20	20
Network	2.5	20
Correction algorithm	10	30
Power Converter Control	20	50
Total	52.5	120

Table I : Summary of time delays for LHC orbit control

5. Controller Design

5.1 Definitions

Before entering the subject of control engineering, some definitions are reviewed here. More information on this subject can be found in literature on control engineering [10,11].

- The system to be controlled is usually referred to as *the plant*.

- An *open loop system* is a system where the measurements from the output (the sensors) are not used to tune the inputs (the actuators). Predefined actuator settings are provided to make the plant behave as desired. No correction of the setting is done afterwards. An example of an open loop system is when a LHC reference magnet is used to modify the corrector magnet settings in the machine.

- *Feed-forward* control is where the actuator settings are changed in accordance to changes in reference or other settings. Even in this context, the correction has no direct effect on the reading of the sensor. An example of a feed forward system is the SPS tune correction : the tune is measured on one cycle and corrected a few cycles later.

- *Feedback* or *closed loop* control is where the actuator has a direct effect on the sensor reading. This way the disturbances entering the system can be compensated for. In many slow stable plants (where the deviations do not accumulate), the operator in fact computes the corrections necessary to bring the plant back to its desired operating point. The subject of control engineering deals with automating this process. The LEP Q-loop was an example of a feedback system.

- The *closed loop bandwidth* of the system plays an important role in disturbance rejection. The closed loop bandwidth is defined as the maximum frequency at which the feedback loop can reduce the effect of sinusoid injected into the loop. For example, a feedback loop with a closed loop bandwidth of 1 Hz will suppress the amplitude of a

sinusoidal error signal at a frequency of 0.5 Hz with a factor 2, at a frequency of 0.1 Hz with a factor 10 and so on.

- The *sampling time* is time between two consecutive observations and corrections. Faster sampling gives better feedback performance because one observes and corrects the actuator on a shorter timescale. If the system is unstable (a small error pushes the system away from its operating point and the error accumulates), then the sampling time needs to be shorter than the natural time scale of the plant, so that the actuator settling can be detected and corrected if needed. However, for stable plants, it is possible to have a sampling time that is much longer than the natural time scale of the plant. The sampling theorem that states that in order to reconstruct a continuous signal from samples of that signal, one needs to sample at least *twice* as fast as the highest frequency of that signal. In practice however, the theoretical lower bound is not sufficient in terms of the quality of the desired time response and sampling rates of *20 to 30 times* the closed loop bandwidth are used. The purpose is to reduce the delay between a command and the system response to this command.

- The *gain margin* of a closed loop system is the factor by which the gain can be increased before the system becomes instable which is usually when the phase shift is approaching 180 degrees. The *phase margin* is the difference between the induced phase shift and 180 degrees when the gain of the system is equal to 1. The phase margin is a measure of how much additional phase lag can be tolerated in the loop before instability occurs.

5.2 Design approach

There are two basic design approaches based on classical and modern control theory. *Classical control theory* is a *frequency domain* approach using the Laplace transform and techniques such as root-locus and frequency response. With techniques based on transfer functions, it is possible to design a classical three-term controller (PID). Modern control theory is a *time domain* approach based on ordinary differential equations. It is also termed the state-space approach, since it uses the concept of states – the minimum set of variables necessary to completely describe the current status of the system. Although the state space method is very different from the transform technique, one ends up with nearly identical controllers in the end.

In addition, there is a principle difference depending on whether time is treated as a continuous or discrete entity. If time is discrete, we enter the area of *digital control* and digital controllers are used by most present day control systems. The frequency domain approach for the digital control theory is uses the z-transform (instead of Laplace transform for continuous systems), while the time domain approach uses difference equations (instead of differential equations for continuous systems).

	Continuous	***Discrete***	***Techniques***	***Literature***
Frequency Domain	Laplace transform	z-transform	Transform techniques	Classical Control
Time Domain	Differential equations	Difference equations	State space	Modern Control

Table II : Overview of standard design approaches in control engineering

In what follows, we will solve the problem of LHC orbit control in the frequency domain (transform techniques) using a discrete setting. A detailed description of this method can be

found in [11]. It will be shown that a simple PI controller is sufficient to place the time constants of the system at their desired locations.

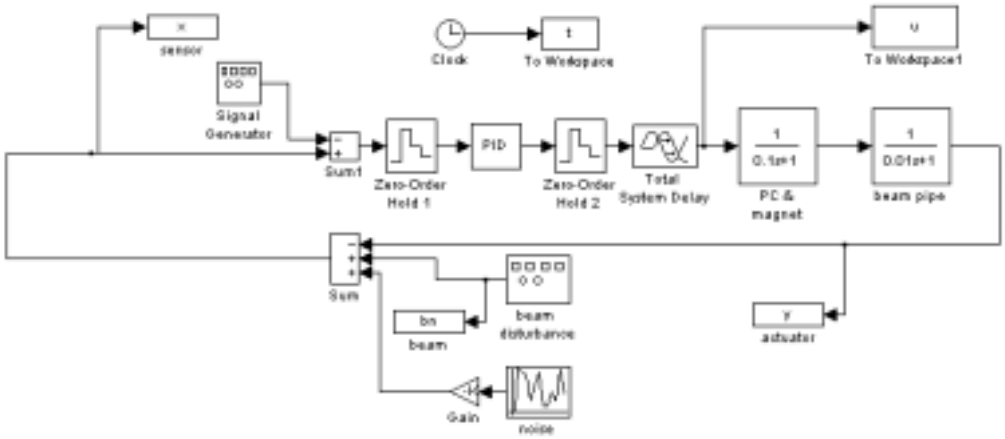

Figure 5 : Matlab Simulink® simulation model for LHC orbit control.

5.3 Design of a Discrete Controller in the frequency domain

Figure 5 shows the Matlab Simulink model that was used to design the controller using proportional (P), integral (I) and derivative (D) control terms. The discrete time transfer function of the PC magnet and the beam pipe has a sampling time of 100ms (10Hz). The time constant of the Power Converter and magnet is 0.1 s (assuming operation at injection energy) and the time constant of the beam pipe is 0.01 s. The system delay is approximately one sampling interval to account for which adds an extra pole at the origin. The discrete overall transfer function is:

$$G(z) = \frac{0.59z + 0.04}{z^2(z - 0.36)} \qquad (5.1)$$

A PI controller can be designed for this transfer function using the standard loop shaping technique of “pole-zero” matching. The PI controller transfer function is given by:

$$H(z) = K_p + \frac{K_i}{z-1} = K_p \frac{z - (1 - K_i/K_p)}{z-1} \qquad (5.2)$$

This PI controller is introducing a pole at *z=1* and leaves a zero to choose. To cancel the pole of the plant with the zero of the controller, it is required that $1\text{-}K_i/K_p = 0.368$ or $K_i = 0.632\ K_p$. This leaves the gain K_p to choose. With a large K_p we get a good disturbance rejection, but we also increase the risk of making the system unstable. Root-locus is the standard technique to deal with this problem. Figure 6 shows the root locus plot for our system indicating the poles (indicated with x) and zeros (indicated with 0) of equation 5.1 using a controller with a gain $K_i = 0.632\ K_p$ and a gain K_p that is varied from zero to infinity.

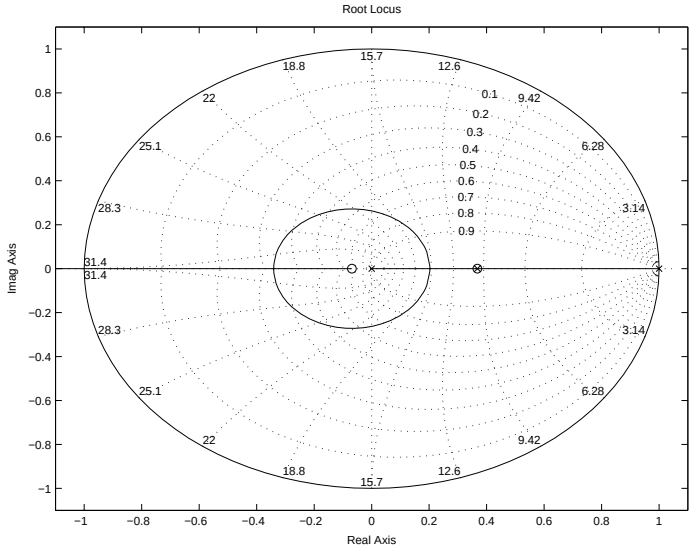

Figure 6 : Root locus plot for the LHC orbit feedback loop using PI control. Poles are indicated by x and zeros by a 0.

On the horizontal axis (from left to right) one finds a zero due to the sampling, a pole at the origin due to the overall time delay, a “matched” pole-zero pair due to the Power Converter & magnet and the PI controller and, finally, a pole from the PI controller.

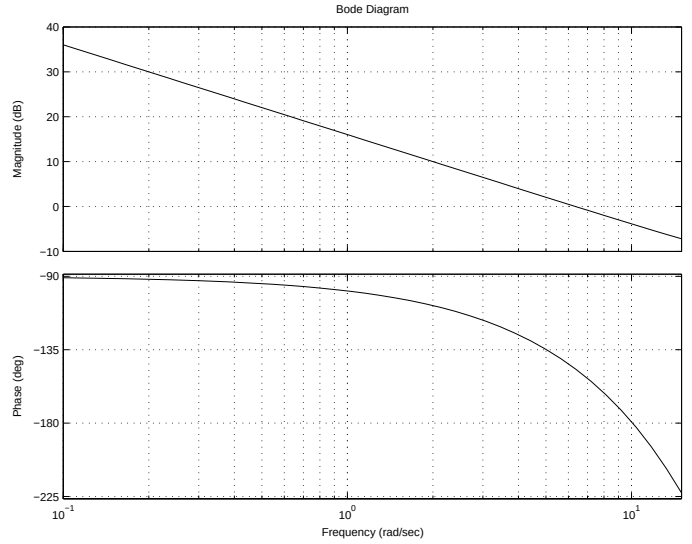

Figure 7 : Bode plots corresponding to the Root Locus of figure 6.

If a proportional gain of 1 is selected, the system can reduce a sinusoid at 0.1 Hz by a factor 8 but there is no reduction for signals at 1 Hz. This can be deduced from the corresponding Bode plot given in figure 7. This figure also indicates the phase margin of 33.2 degrees and the gain margin of 1.6.

Other controllers and/or optimisation procedures have been considered and are listed in table III. For example, instead of fixing the zero of the controller to coincide with the pole of the plant, it can be considered as an additional decision variable to improve the gain at 1 Hz. An optimisation routine allowed choosing the zero while enforcing a phase margin of 30 degrees. The optimised PI controller shows an improvement in gain at 1Hz at the cost of reducing the gain at 0.1 Hz.

Alternatively, a PID controller can be designed to increase the bandwidth and reduce the influence at 1Hz. Unfortunately, only marginal improvement (30% reduction for 1 Hz) was

obtained (table III) which is basically due to sampling at 10 Hz while having a time delay of 100ms. This means that there is not much margin for disturbance rejection at 1Hz. It should also be noted that the derivative term makes the system more sensitive to BPM noise, which could pose a problem with low intensity (pilot) bunches.

To compensate the time delay, a Smith predictor can be used. A Smith compensator is feeding back a simulated plant output to cancel the true plant output and then adding in a simulated plant output without the delay. The derivative term does not bring an improvement in performance anymore, so only a PI controller with Smith predictor has been considered (table III). It can be seen that the gains are considerably higher with the Smith predictor. However, Smith predictors are more difficult to tune than simple PI controllers.

A sampling frequency of 20Hz gives a consistently better performance in any case.

Controller	Gain at 1 Hz	Gain at 0.1 Hz	Kp	Ki	Kd
PI – 10 Hz (Standard)	1.01	10.56	1	0.632	-
PI – 10 Hz (Optimized)	1.24	7.92	1.49	0.493	-
PID – 10 Hz (Optimized)	1.32	15.86	1.39	1	0.605
PI – 10 Hz (Smith Predictor)	2.2	22.2	2.2	1.39	-
PI – 20 Hz (Standard)	1.25	12.50	1	0.394	-
PI – 20 Hz (Optimized)	1.39	10.45	1.45	0.327	-
PID – 20 Hz (Optimized)	1.80	23.4	1.42	0.739	1.65
PI – 20 Hz (Smith Predictor)	6.2	62.6	5	1.74	-

Table III : Performance comparison of different controllers.

6. Conclusions

Orbit control will play an important role in the LHC due to the tight aperture. During a physics run, the closed orbit will be perturbed by the movement of the triplet quadrupoles, the β-squeeze and by the decay of the persistent currents. The maximum frequency of the orbit perturbation has been estimated at 0.05 Hz. In this paper, it has been shown that a properly designed feedback loop can eliminate such orbit distortions during injection and at the beginning of the ramp. For example, a proportional–integral control loop that samples at 10 Hz and that has a total time delay of 100 ms can reduce orbit errors at frequency of 0.1 Hz with a gain of 8. In all cases, a considerably better performance can be obtained when the sampling frequency is increased and/or when the total time delay is reduced.

Acknowledgements

The authors would like to thank F. Bordry, A. Burns , A. Butterworth, L. Hendrickson, W.Herr, T. Himmel, M. Lamont, R. Lauckner, J. Pett, P. Ribeiro and L. Vos for their contributions.

References

[1] J. Wenninger *et al*., *Emmittance optimisation with dispersion free steering at LEP*, CERN-SL-2000-078-OP.

[2] R. Brinkmann, J. Rossbach, *Observation of closed orbit drift at HERA covering 8 decades of frequency*, Nucl. Instr. Meth. A 350 (1994) 8-12.

[2] L. Vos, A. Verdier, Ground Motion Model for the LHC, *22nd Advanced ICFA Beam Dynamics Workshop on Ground Motion, Stanford, USA, 6 - 9 Nov 2000*.

[3] R. Wolf *et al., The Expected Persistent Current Field Errors in the LHC Main Dipole and Quadrupole, LHC-PROJECT-NOTE-230, August 2000.*

[4] T. Wijnands *et al.*, *Requirements for the real time correction of decay and snap back in the LHC superconducting magnets, LHC-PROJECT-NOTE-221, April 2000.*

[5] W. Herr *et al*., *http://wwwslap.cern.ch/collective/zwe/lhcbb/Welcome.html*

[6] W. Herr *et al*., *http://wwwslap.cern.ch/~zwe/cocu/cocu.html*

[7] W.H Press *et al*, Numerical Recipes in C, Cambridge University Press, 2nd Edition

[8] F. Bordry, Power Converters, *Proceedings of the XIth Chamonix workshop, 15 –19 January 2001*.

[9] O. Bruening, Accumulation and Ramping, *Proceedings of the Xth Chamonix workshop, 17-21 January 2000.*

[10] L. Vos, Loose thoughts on LHC feedback control, *Private Communication.*

[11] T. Wijnands *et al*, Supporting QoS for LHC controls applications using ATM and LynxOS, *LHC-project report 231, May 2000*.

[12] G. Franklin, J. Powell, A. Emami-Naeini, Feedback Control of Dynamic Systems, *Addison Wesley Longman Inc., Menlo Park U.S.A., 3rd edition, 1994.*

[13] G. Franklin, J. Powell, M. Workman, Digital Control of Dynamic Systems, *Addison Wesley Longman Inc., Menlo Park U.S.A., 3rd edition, 1998.*