COMMISSIONING AND INITIAL PERFORMANCE OF THE LHC BEAM-BASED FEEDBACK SYSTEMS

R. J. Steinhagen, A. Boccardi, A. Butterworth, E. Calvo Giraldo, R. Denz, M. Gasior, J.L. Gonzalez, S. Jackson, L. K. Jensen, O. R. Jones, Q. King, M. Lamont, G. Kruk, S. Page, J. Wenninger, CERN, Geneva, Switzerland

Abstract

The LHC deploys a comprehensive suite of beam-based feedbacks for safe and reliable machine operation. This contribution summarises the commissioning and early results of the LHC feedback control systems on orbit, tune, chromaticity, and energy. Their performance – strongly linked to the associated beam instrumentation, external beam perturbation sources and optics uncertainties – is evaluated and compared with the initial feedback design assumptions.

INTRODUCTION

To minimise cross-talk between the individual feedback (FB) systems, the control of orbit, tune (Q), chromaticity (Q'), coupling ($|C^-|$) and the RF radial loop for energy compensation are combined to one single feedback controller [1]. This controller receives, reconstructs and checks the measurement of about 2200 inputs from the beam position monitors and the Q/Q' diagnostics systems, calculates derivatives such as Q' and dispersion orbit components, and computes – based on a singular-value-decomposition (SVD) based approach – the magnetic field and RF frequency corrections that are necessary to stabilise the beam around a given beam parameter reference. These corrections are sent to about 1300 corrector circuits and the RF cavity systems, each controlled by a function generator (FGC) [2, 3].

Until now, the LHC beam stability required only slow and small orbit and Q/Q' parameter corrections [4]. Thus an effective FB bandwidth of about 1 Hz and sampling frequency of 25 Hz proved to be sufficient. Due to the large number of active elements involved, identification of errors and failures as well as the deployment of mitigation procedures is a key-aspect of the FB since an undetected transients could lead to loop instability, particle loss and eventually even to a beam dump for protection.

The controller input data is re-published and used by user-level applications for orbit and Q/Q' monitoring and correction, slow software-based interlocks, logging and interaction point optimisations.

INITIAL FEEDBACK OPERATION

The orbit, Q and Q' diagnostics and Q/Q' feedback systems – based on the Base-Band-Tune (BBQ) measurement system – were generally considered to be a 'work horse' from Day-I of LHC commissioning [4]. Prior to first circulating beams and energy ramps, the targeted initial commissioning sequence was: orbit (OFB), energy (radial loop), Q' and later Q- and $|C^-|$- feedback. However, in response to large portions or the entire beam being lost due to large tune drifts during the first ramps, the commissioning of the Q-FB was given priority and was thus operated early on.

In order to isolate potential problems out of the more than 3400 elements involved, the initial planning foresaw a thorough test of each of the individual feedback sub-components and associated system responses. This was eventually limited to checks of the orbit, tune and sextupole corrector circuits and open-loop feedback polarities. Nevertheless, the open- and closed-loop responses proved to be stable and well in agreement with the model transfer functions. Figs. 1(a) and 1(b) show the corresponding orbit and tune feedback responses to external perturbation sources.

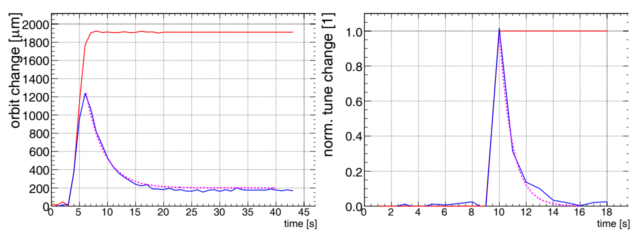

(a) Orbit-FB Step-Response (b) Q-FB Step-Response

Figure 1: Measured (blue) and model (dashed, magenta) Orbit- (a) and Q-FB (b) response to an external perturbation (red): $\Delta x = 2\,\mathrm{mm}$ and $\Delta Q = 0.003$.

The FB responses have been tested using an independent input to the corrector circuits, in case of the orbit using a single three-corrector bump. The residual steady-state error in the OFB response is caused by the chosen locality of the SVD algorithm that prevented the correction of local bumps but which can easily be adjusted. The closed-loop responses were essentially limited only by the quality of the associated beam instrumentation, in good agreement with the model assumption and confirmed the absence of polarity and large calibration errors.

While the initial effective bandwidth was configured to 0.1 Hz, the absence of large transfer function errors should easily allow an effective bandwidth of 1 Hz. Higher bandwidths may require a more detailed assessment of all involved delays, corrector and beam transfer functions.

ORBIT-FB PERFORMANCE

The orbit stability without feedback during the ramps was typically in the order of 1 mm r.m.s. and required only a few feed-forward (FF) corrections based on the monitoring of previous ramps. The actually orbit stability at collision energies was essentially determined by slow tidal variations in the order of $100\,\mu\mathrm{m}$ only. While the Orbit-FB was available it was only run during a few test ramps as the residual observed stability was sufficient and rarely caused orbit-drift induced losses. Fig. 2 shows an example of orbit r.m.s. evolution with and without FB for two successive ramps. In this particularly case, the Orbit-FB could maintain an orbit stability of better than $70\,\mu m$ compared to orbit perturbations of up to about $350\,\mu m$ r.m.s. without Orbit-FB. While the predicted worst-case orbit drift estimates ranged up to $25\,\mu\mathrm{m/s}$ – in particular during the β^*-squeeze – the actual measured perturbations were usually quite small in the order of about $1\,\mu\mathrm{m/s}$. This greatly relaxes the requirement on the effective loop bandwidth. As for the step response shown in Fig. 1(a), the residual orbit r.m.s. with FB 'on' is limited by the chosen locality of the correction algorithm and can be easily reduced to a few microns assuming the absence of systematic BPM errors. Beside the operational integration into day-to-day operation, one of the main functionality that remains to be commissioned is the semi-automated controller adjustments to accommodate the changing optics during the β^*-squeeze in the interaction points.

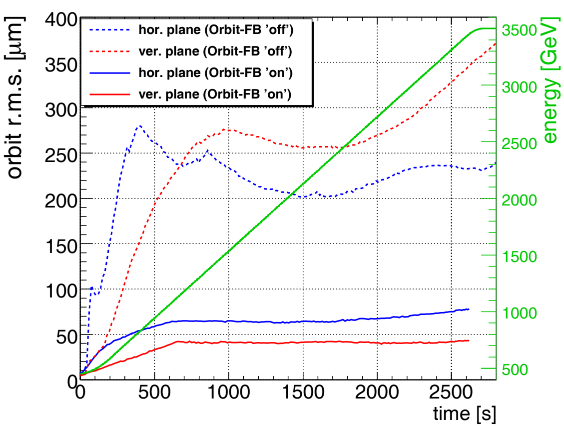

Figure 2: Evolution of the r.m.s. orbit deviation during the ramp with (solid) and without (dashed line) feedback.

The main performance limitation of the OFB is linked to the systematic BPM dependence on temperature and bunch intensity that causes measurement errors on the orbit measurement of up to $300\,\mu\mathrm{m}$. As described in [7], temperature changes in the acquisition electronics generate measurement drifts in the order of $100\,\mu\mathrm{m}/^{o}\mathrm{C}$. This temperature effect is being addressed by a local temperature control of the BPM front-ends. During the recent commissioning with high bunch intensities it was found that the systematic dependence on beam intensity causes systematic orbit measurement errors of up to $300\,\mu m$ for bunch intensity changes from a about $1.1 \cdot 10^{11}$ down to about $5 \cdot 10^{10}$ charges per bunch. While these drifts are acceptable during the ongoing initial commissioning with low beam intensities, further measurements with beam are required to fully assess the possibility of calibrating and compensating for this effect in view of the much tighter orbit control requirements with high-intensity beams.

Alternate SVD Orbit Correction

Nearly all Orbit-FB rely on an SVD-based approach, using the orbit response matrix (ORM) $\mathbf{R}$ to compute the change of orbit corrector kicks $\Delta\vec{\delta}$ in response to measured orbit deviations $\Delta\vec{x}$:

$$\Delta\vec{x} = \mathbf{R}\cdot\Delta\vec{\delta} \;\text{ with }\; \mathbf{R} = \mathbf{U}\cdot\lambda\cdot\mathbf{V}^T, \tag{1}$$

with $\mathbf{U}$ being a rectangular, $\mathbf{V}$ a square and λ a diagonal matrix containing the eigenvalues of $\mathbf{R}$. The individual vectors $(\mathbf{U})_i$ and $(\mathbf{V})_i$ span two orthonormal bases in orbit- and corrector-space. Typically, the inverse problem is solved by creation of a 'pseudo-inverse ORM'

$$\tilde{\mathbf{R}}^{-1} = \mathbf{V}\cdot\tilde{\lambda}^{-1}\cdot\mathbf{U}^T \tag{2}$$

with $\tilde{\lambda}^{-1}$ containing the inverted eigenvalues of λ for non-singular values and zeros otherwise. Lower eigenvalues typically correspond to global orbit patterns, whereas higher eigenvalues to local orbit bumps. Often, the near-singular values are deliberately set to zero to adjust the trade-off between the precision of the orbit correction and sensitivity to measurement noise or failures.

Fig. 3 illustrates the trade-off between orbit attenuation and sensitivity to measurement noise as a function of the used number of eigenvalues $\#\lambda_{SVD}$. For increasing $\#\lambda_{SVD}$ the precision of the orbit correction improves (which is needed in regions around the collimators and interaction points) but the corrections also become more prone to noise and failures of individual BPMs.

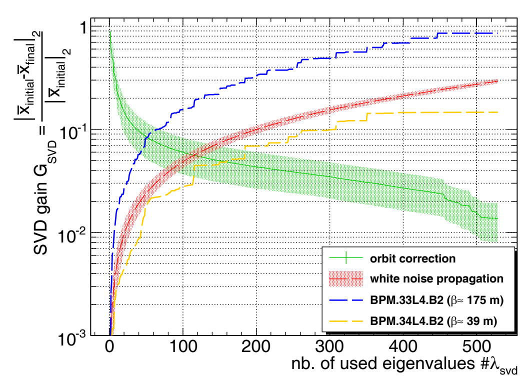

Figure 3: Dependence of the orbit correction gain and sensitivity to random ('white') noise, failures of a single BPM at a maximum and minimum β-function on $\#\lambda_{SVD}$.

The Orbit-FB implements also an alternate correction algorithm that mitigates the compromise between having local orbit control requiring large $\#\lambda_{SVD}$ and increasing the sensitivity to measurement noise by decomposing the matrix-vector multiplication of Eq. 2 with the orbit shift $\Delta\vec{x}$ into eigen-solutions

$$\Delta\vec{\delta}(t) = \sum_{i=0}^{n} D_i\left(\frac{a_i}{\lambda_i}(\mathbf{V})_i, t\right) \text{ with } a_i = (\mathbf{U})_i^T\cdot\Delta\vec{x} \tag{3}$$

and to treat the each deflection vector of the given eigenvalue by an independent feedback controller function $D_i(x,t)$. For $D_i(x,t) = x$ this scheme is identical and has about twice the numerical complexity compared to the classic matrix-vector multiplication. Using this scheme it is possible to use all eigenvalues, thus to correct for very localised perturbations, while – in view of long-term orbit stability – mitigating the effect of measurement noise by reducing the bandwidth of the corresponding higher eigenvalues. Alternatively, a fast local control can be provided while reducing the bandwidth for global-type orbit correction. Also, fast changes of $\#\lambda_{SVD}$ are possible for which the classic approach requires a re-computation of $\tilde{\mathbf{R}}^{-1}$ that could take several seconds.

TUNE-FB PERFORMANCE

Due to the high sensitivity of the BBQ and residual tune oscillations, most day-to-day Q diagnostics and feedback operation could be based on a simple Fourier-based analysis of the BBQ spectra. With the Q-FB being switched 'on', the actual tune stability was usually kept better than 10^{-3} around the nominal tunes ($q_h = 0.28$ & $q_v = 0.31$) during injection and better than 10^{-4} during the ramp [4]. Fig. 4(a) shows the evolution of the bare tunes – the tunes that would have been obtained if the FB would have been switched 'off' – during two successive ramps for an ideal operational scenario with the same magnetic history of the main dipole and quadrupole circuits prior to both ramps and with no other additional correction. While the overall

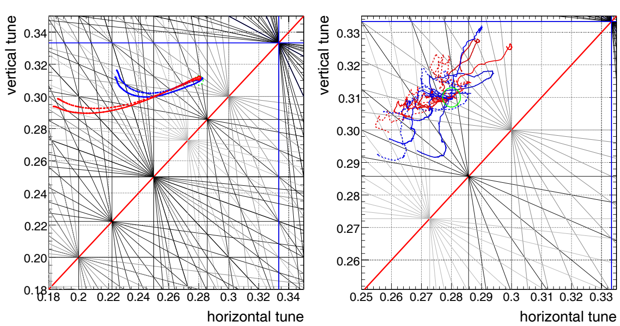

(a) FB-piloted ramps without FF (b) FB-assisted ramps with FF

Figure 4: Reconstructed Beam 1 (blue) and Beam 2 (red) tune drifts during the first ramps. The variations between ideal (left) and de-facto pre-cycle scheme (right) is visible. Resonance lines are indicated up to the sixth order. The actual tune stability was better than 10^{-3} in all shown cases.

recorded trims are large and would certainly have caused multiple tune resonance crossings and thus beam losses in the absence of the Q-FB, the difference between the individual ramps are small and correspond to a fill-to-fill reproducibility of about 3 to $5 \cdot 10^{-3}$. Hence, these FB corrections were incorporated into the successive ramps as part of the beam-based FF scheme to reduce the overall required tune trims and dependence on the Q-FB. However, for later ramps the achieved stability varied significantly and reached variation as much as ± 0.03 as illustrated in Fig. 4(b) showing the bare tunes for 10 ramps during a one month period. Large parts of these non-reproducibilities are assumed to be due to strong variations of the magnetic pre-cycle history of the machine. In most of these cases (e.g. after an access or trip of a single sector) the default pre-cycle was substitutes by a faster version in favour of minimising the beam down-time to allow faster commissioning progress [5, 6].

The residual short-term stability at injection of typically about $\pm 5 \cdot 10^{-4}$ was found to be random noise, and since it scaled down with energy is assumed to be due to current ripple in strong individually powered quadrupole magnets such as the warm insertion or triplet quadrupoles, or possibly tune trim quadrupoles as they are usually used with very small mA-level currents, while their maximum current rating is 600 A. Further studies are required to fully assess and mitigate these effects. In any case, since the stability scales with energy and improves by about a factor of eight, these effects are less important at top energy.

An important limit on the Q-, Q'- and $|C^-|$-FB operation is given by the magnet quench-protection-system (QPS) that issued false-positive quench- and thus energy extraction triggers if the real-time FB trims exceeded the maximum accepted QPS current acceleration rates. The tune corrector circuits' QPS has recently been modified in response to this and now filters the FB specific real-time trim induced accelerations. Similar mitigations are being investigates for the skew-quadrupoles and sextupole circuits.

Due to these limitations the Q'- and $|C^-|$-FB loops have not yet been closed during the ramp or at high energies. However their derived measurements been used to improve the Q' and coupling evolution during the ramp and β^*-squeezes as part of the fill-to-fill feed-forward corrections.

CONCLUSIONS

The commissioning of the beam-based feedbacks and associated diagnostic chains advanced well during the first days with beam and facilitated early-on a fast and reliable LHC operation. In response to tune drift related particle losses during the first ramps, the commissioning of the Q-FB was given priority and and achieved tune stabilities of a few 10^{-4} at injection down to 10^{-4} for energies above 0.8 TeV. The orbit feedback achieved stabilities of 70 μm r.m.s. during the first ramps and is currently limited by the chosen locality of the correction. In the long-term however, BPM dependencies on crate temperature and bunch intensity leading to drifts of up to 300 μm will have to be addressed.

REFERENCES

[1] A. Boccardi et al., CERN-BE-2009-002, 2009
[2] M. Bastos et al., LHC Project Report 1044, 2007
[3] F. Bordry, H. Thiesen, LHC-Project-Report-258, 1998
[4] R. J. Steinhagen et al., BIW'10, 2010
[5] M. Lamont et al. "Q/Q' Reproducibility During the Early LHC ramps", private communication
[6] E. Todesco, "The magnetic model of the LHC in the early phase of beam commissioning", these proceedings
[7] R. Jones, LHC-Performance-Note-006, 2009