REAL-TIME BEAM CONTROL AT THE LHC

R. J. Steinhagen, CERN, Geneva, Switzerland

Abstract

At the LHC, real-time feedback systems continually control the orbit, tune, coupling, and chromaticity. Reliable and precise control of these parameters is essential for a safe and reliable machine operation. This contribution summarises the feedback performance during LHC's first full year of operation.

INTRODUCTION

The requirements of LHC's key beam parameters strongly depend on the capability to control particle loss inside the accelerator. Not only driven by machine protection, collimation and quench prevention, but also commissioning and operational efficiency, the function of these systems depends critically on the stability of orbit, energy, tune (Q), chromaticity (Q') and betatron coupling (C^-), and imposes significant constraints on the maximum allowed beam excursions, traditionally required to measure Q and Q' to a few μm. During the LHC re-start at the end of 2009 and through early 2010, the orbit, Q and Q' diagnostics and feedback systems – based on the Base-Band-Tune (BBQ) measurement system – were generally considered to be 'workhorses' and facilitated a fast and reliable commissioning [1, 2, 3].

Guided by the expected perturbation and tight requirements, Q' was initially considered to be the most critical parameter, defining lifetime and dynamic aperture of the beam, followed by C^- especially during the start of ramp, enabling the control of the other beam parameters. Prior to first circulating beams and energy ramps and led by what was easiest to commission, the targeted sequence was: orbit, energy (radial loop), Q' and at a later stage Tune- and Coupling-Feedback. However, in response to large portions or the entire beam being lost due to large tune drifts during the initial ramps, the commissioning of the Tune-FB followed by the Orbit-FB were given priority and were thus operated early on and during almost every fill.

Q/Q'-PERFORMANCE

While the Tune Phase-Locked-Loop (PLL) has been commissioned and used during some ramps, due to the BBQ's nm-level sensitivity, most day-to-day Q/Q' diagnostics were nevertheless performed based on passive monitoring of the beam spectra only, limiting potential impact on beam size growth. The change of paradigm of deriving the Q/Q' from passive monitoring instead of resonant excitation of the beam required some adaptations in the digital post-processing, which after the appropriate strategy was established performed surprisingly well and soon became the base-line mode of operation of the feedbacks.

Tune-FB

The Tune-FB performance was steady over the year and largely dominated by the snap-back at the start of the ramp as shown in Figure 1, showing the superimposed residual tune stability for both beams.

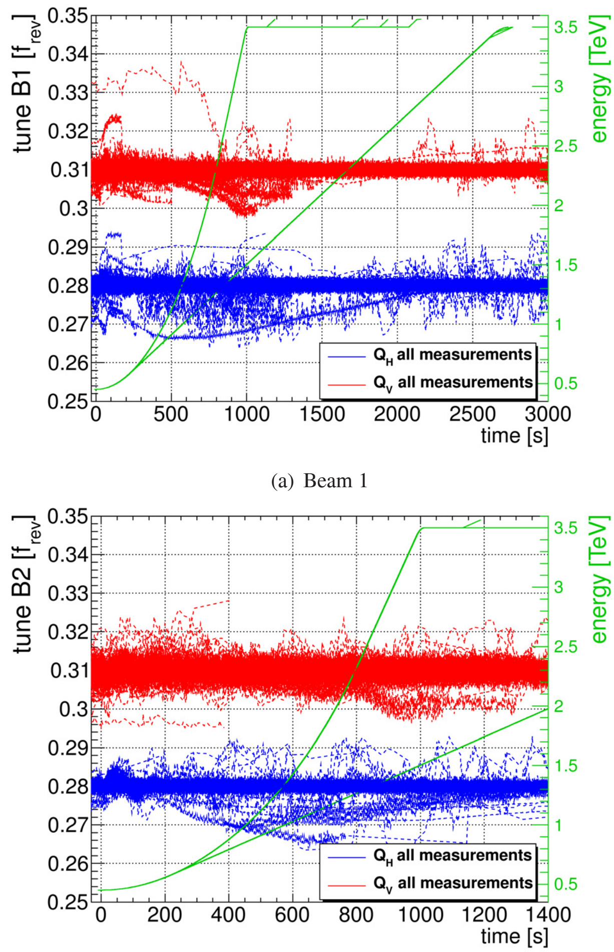

(a) Beam 1

(b) Beam 2

Figure 1: Residual tune stability. Outliers are due to a few test ramps without Q/Q' feedbacks for diagnostics purposes and temporary BBQ unavailability.

Initially, very conservative feedback settings were chosen, which resulted in exceeding the required nominal tune stability by about 10^{-2} mainly during the first 120 seconds of the ramp. At a later stage, once operating the LHC with ions and after a reliable BBQ and feedback operation was widely affirmed, this stability was further improved to below $3 \cdot 10^{-3}$. The presently achieved stability is limited rather by the resolution, stability and reliability of the Q/Q' diagnostics rather than the feedback controller or loop itself.

Experience at RHIC has shown that a reliable control of Q implies also a reliable control of coupling which has been adapted also for the LHC[4]. However, coupling proved to be less of an issue, was very reproducible and kept below the $|C^-| < 3 \cdot 10^{-3}$ level, once the fill-to-fill recurring quadrupole alignment driven perturbations were compensated for. In addition, since the orbit was kept constant in the arc sextupoles, feed-down effects could be kept to a minimum.

Being used during nearly every ramp and squeeze to physics, losses could be kept at a minimum. Out of a total of 275 ramps, excluding the early ramps in 2009, a total of 155 (122) ramps achieved more than 99%, 169 (155) ramps more than 98% and 178 (168) ramps more than 97% transmission for B1 (B2). Only 12 (10) ramps were lost with some feedback involvement, out of which 6 (5) were during the initial 3.5 TeV commissioning.

Chromaticity-FB

The expected large chromaticity perturbation due to the known $b3$-field component of the main dipoles has been corrected using the sextupolar corrector spool pieces foreseen for this purpose. Initially, a first-order model based correction was applied relying on the magnetic field measurements that have been performed for all LHC dipole magnets before their installation. Non-linear field contributions and dynamic aperture limitations proved to be very small with negligible impact on very low-intensity beams ($n_b \approx 10^{10} protons/bunch$), as seen during the very first acceleration ramps with periods of both large positive ($Q' > 25$) and negative chromaticities ($Q' < -15$) that did not cause excessive transmission losses or beam instabilities. The $Q'(t)$ data taken with the Chroma-FB during these ramps has been used to improve the feed-forward model and incorporated into the next fill's feed-forward correction. The evolution of $Q'(t)$ was measured for a series of consecutive fills and indicated a reproducible behaviour as shown in Figure 2.

The remaining largest fill-to-fill variations occurred as expected during the first 200 seconds of the ramp ('snap-back') reaching up to $\Delta Q' \approx \pm 5$ only. Once reaching 3.5 TeV another decay of about 6 units of chromaticity is visible. To allow this decay to settle, the ramp was artificially extended by about 6 minutes to compensate for this effect. In between, the chromaticity was found to be stable within about $\Delta Q' \approx \pm 2$ which indicates that, besides the snap-back, most of these effects are well compensated by LHC's feed-forward function alone and nearly down to nominal requirements. The very good reproducibility is enforced by strict pre-cycling of all magnets following physics, accesses, circuits being 'off' and other irregularities. The continuing $Q'(t)$ measurements during injection and ramps are used to study and improve the analytic field description model for the LHC dealing with the dynamic decay and snap-back compensation (FiDeL, [5]).

ORBIT-FEEDBACK

Since the beginning of 2010, the Tune-FB has been routinely complemented by the Orbit-FB during ramp and squeeze[1]. The present Orbit-FB correction relies on standard approach of using the pseudo-inverse of the orbit-response-matrix (ORM) that is computed through Singular-Value-Decomposition (SVD) of the ORM, followed by an optimal controller which converts the dipole corrector deflections into currents. During the first fills, the near-singular solutions were truncated (TSVD) by setting the inverse of near-singular eigenvalues to zero. Often, more eigenvalues were deliberately removed to make the correction less sensitive to BPM noise but also allowing local orbit bumps to creep in. Since early 2010, the more common regularised SVD approach is used, where the pseudo-inverse of the eigenvalues is given by

$$\lambda_i^{-1} := \frac{\lambda_i}{\lambda_i^2 + \mu} \tag{1}$$

with $\mu > 0$ being the regularisation parameter (aka. Tikhonov regularisation). This scheme proved to be more robust with respect to optics errors and allowed re-using the same pseudo-inverse ORM for injection, ramp and the various individual squeeze steps, instead of being forced to change a couple of dozen optics as was initially foreseen. As a beneficial side-effect: while using more eigenvalues also implies a higher sensitivity to localised BPM errors and noise for the TSVD-based approach, the regularisation mitigates this effect by effectively reducing the bandwidth for local orbit corrections – thus implicitly increasing the averaging of the BPM reading – while maintaining a fast feedback response for global corrections that rely on the reading of many BPMs [3].

Using this scheme the Orbit-FB could maintain orbit stabilities of typically better than 70 μm globally and 20 μm in the arcs compared to orbit perturbations of up to about 1 mm without Orbit-FB. Most of the remaining orbit perturbations are due to programmed dynamic reference changes around the experimental insertions during ramp and squeeze. The main orbit perturbations at collision energies is driven by slow tidal variations in the order of $\pm 100\,\mu$m and with typical periods of about 12 hours.

The main performance limitation of the Orbit-FB is linked to the systematic BPM dependence on temperature and bunch intensity that initially caused errors on orbit measurement of up to 300 μm. As described in [7, 8], temperature changes in the acquisition electronics generate measurement drifts in the order of $100\,\mu\mathrm{m}/^{\circ}\mathrm{C}$. For the time being, this effect is suppressed by data post-processing to below $100\,\mu$m using the crate temperature as an estimate of the electronics card temperature. A full temperature control of the front-ends is under investigation.

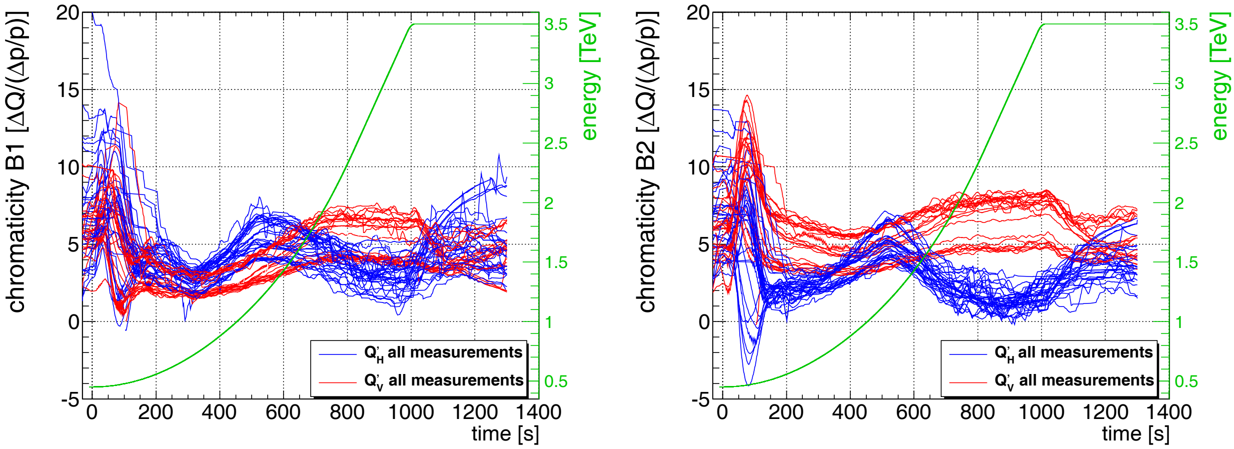

(a) Beam 1 (b) Beam 2

Figure 2: Residual superimposed $Q'(t)$ stability during successive ramps.

FEEDBACK ISSUES AND MITIGATION

The few beam dumps involving feedbacks were limited to the initial setup and commissioning during the first months and had a small (below percent-level) impact on overall machine operation [10, 11]. Most of the beam dumps where feedbacks were involved were due to either false-positive QPS trips which have been mitigated by introducing a dead-time in the evaluation of the QPS threshold, and due to locking of the BBQ tune diagnostics on non-tune resonance lines in the spectrum. The tune tracker was modified early on in response to this, and most of these non-tune interference lines have now been identified and eliminated using a multi-stage, median-filter based search algorithm that removes lines based on their bandwidth.

Working with Residual Tune Oscillations

The initial tune diagnostics design assumed no residual tune signatures on the beam and hence a constant driving of the beam (e.g. a 'kick', 'white noise', 'chirp' or 'PLL') was envisaged. To limit the required excitation levels and consequently minimising the resulting potential emittance blow-up, the highly-sensitive BBQ system was developed, which has been further exploited by a real-time FFT spectrum analysis and PLL system[2]. The working hypothesis was that the BBQ's nm-level sensitivity would be sufficient to operate below the oscillation level, which would/could be damped by the ADT, and which would impact machine operation or protection. Initial tests at the RHIC, SPS and Tevatron, and likewise early experience after the start-up and present LHC operation seemed to confirm this hypothesis with beam: the BBQ can provide a turn-by-turn resolution of better than 30 nm, more than 50 times' sensitivity than any other LHC systems (ADT: 1 μm [15], BPM: 50 μm [7, 8]). At the same time, ever-present residual tune oscillations are visible on the LHC beam with amplitude in the order of 100 nm to a few micro-metres. This "luxurious" 30 to 40 dB signal-to-noise ratio facilitated a passive monitoring, tracking and feedback without additional excitation, which proved to be sufficiently reliable from day one, controlling large tune variations during almost every LHC ramp (and most squeezes). The substantial resolution also helped to identify other beam perturbation issues such as electromagnetic interferences, the 'hump', and other effects documented in [16].

While these μm-level oscillations are a-priori beneficial for a passive detection of the tune, they are incoherent 'noise' from a FFT or PLL diagnostics point of view. Regardless of whether using a driven FFT- or PLL-based diagnostic tune system, the beam needs to be excited about 20-30 dB above this 'noise' to recover the same reliable performance as using residual oscillations only. The corresponding absolute amplitude of about $10 - 100\,\mu$m that is excited on top of the residual tune oscillations are in conflict with collimator requirements ($< 200\,\mu$m and shown to cause beam losses in the machine. Thus driving the beam to such ample signals seemed to be inefficient and less robust compared to the performance achieved with the passive-only system and was considered to be used mainly if the signal dropped.

ADT Interferences

An important issue affecting the reliability and function of the Q/Q' diagnostics and feedback systems is the intrinsically competing requirement of the transverse bunch-by-bunch feedback system (ADT) targeting the minimisation of beam oscillations on the tune frequency and the fact that a certain amount of these oscillations are required to actually measure and stabilise the tune. The nature of these opposing requirements were already recognised during the LHC design phase[13].

The ADT is successfully operated since July, damping injection oscillations on a regular basis, and being kept 'on' also during ramp and collisions with a of damping times of few hundred down to 50 turns[15]. As an intrinsic limiting factor of any feedback, an imporant part of the ADT measurement noise is propagated onto the beam in high-gain operation, as illustrated in Figure 3. This additional beam noise increases the measured BBQ beam noise floor by up to 30 dB and thus reduces the effected tune resolution that is available for the Tune-FB

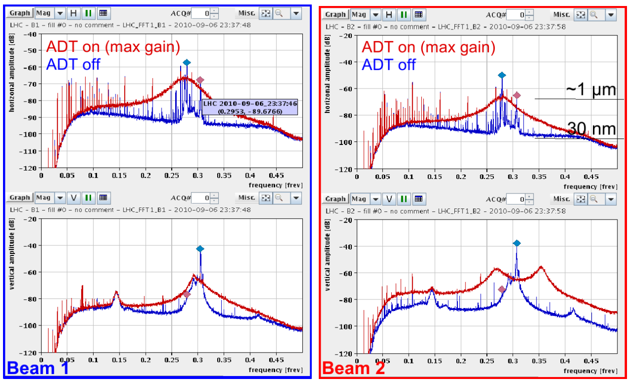

(a) Beam 1 Spectra (b) Beam 2 Spectra

Figure 3: Comparions of BBQ tune spectra with ADT feedback being active with nominal settings (red) and being 'off' (blue). The increase of the beam noise floor and additional introduced structures is visible.

By comparison of the unperturbed and damped spectra, the particular shape of the noise shows a limited relationship between the measured spectra maxima and actual tune-resonance, which initially hindered, and in some cases, prevented reliable operation of the Q/Q'-diagnostics and related feedbacks. For the time being, coupled-bunch instabilities are not yet dominating LHC beam stability and thus the ADT can be operated with reduced gains whenever a precise Q/Q' diagnostics or Tune-FB are required. While this mode of operation is presently acceptable operating with only up to 200 nominal bunches, this may change with increased number and decreasing bunch spacing that is necessary to reach the nominal LHC performance. Thus, investigations are being made to assess the quality and robustness of the tune oscillations that can be recovered from the ADT's actuator signal.

Operational Dependence on Feedbacks

In some cases, the feedback also compensated for effects that were introduced either directly (human and/or feed-forward incorporation errors) or indirectly through feed-down effects that were otherwise not accounted for by day-to-day operations (such as incomplete pre-cycles after accesses, newly measured $Q'(t)$ incorporation into the ramp functions etc.). These examples nicely demonstrated that – even with perfect feed-forward incorporation of the recurring real-time actions – feedbacks can and did provide some additional safety margins to operation by indifferently suppressing and absorbing unexpected perturbations. At the same time, it should be pointed out that without feedback support, the beams would have been probably lost in these cases, which reduces the merit of 'additional' to 'mandatory' safety by the feedbacks. Unfolding the effect of the real-time trims on the tune, out of 275 ramps that were executed in 2010: 56 (83) would have been lost on low-order resonances (3rd, 4th, C-), 150 (157) would have exceeded a $\Delta Q = \pm 0.01$ tolerance, which probably would have caused transmission losses and all were above the $\Delta Q = \pm 0.001$ stability requirement for nominal beams [13]. In order to reduce this dependence on feedbacks, systematic monitoring and transferring of recurring real-time feedback actions into the ramp and squeeze functions are performed more regularly now.

Mains Harmonics

A limited correlation between residual tune stability and intensity transmission during the ramp has been noted, with an initially surprising exception that for fills with stabilities better than 0.005, more intensity was lost than for those with poorer tune control. This is a bit counter-intuitive and would naively suggest not to control the tune. Revisiting the spectra of the given ramps revealed that in these cases the tunes were kept on the horizontal nominal LHC tune working point, which is located exactly on one of the mains harmonic as shown in Figure 4. A set of mains harmonic is visible and more pronounced for high-intensity beams as the BBQ detector becomes more sensitive down to the nm-level. These mains harmonic are typically very small and compatible with the measured and specified main dipole ripple[14]. Their impact is a priori not a big issue and can be easily mitigated by shifting the nominal working points by 0.001 only. Studies for a slightly shifted and completely new working point are under preparation.

PLANNED FEEDBACK MODIFICATIONS

An automatic feedback gain scheduling is planned for 2011, in order to allow a more fine-grained control of the various feedback bandwidths, depending on the operational condition: fast feedback action (/high bandwidth) when fast perturbations are expected (e.g. during the start of the ramp) and slow feedback action (/small bandwidth), which reduces the inevitable noise that is propagated from the beam instrumentation to the beam via the feedbacks. The target is to make the dynamic change dependent on the variation of the residual feedback error signal, but a simple switch will be put in place that will allow the 'high' and 'low' extremes of bandwidth in advance.

CONCLUSIONS

The beam-based feedbacks on orbit, tune and chromaticity performed well in 2010 and facilitated fast and reliable re-commissioning with minimal losses and with near nominal beam parameter stabilities. Analysis of the feedback actions of more than 280 logged ramps indicated that more than half of all fills would have been probably lost without feedback support and the others likely affected by particle

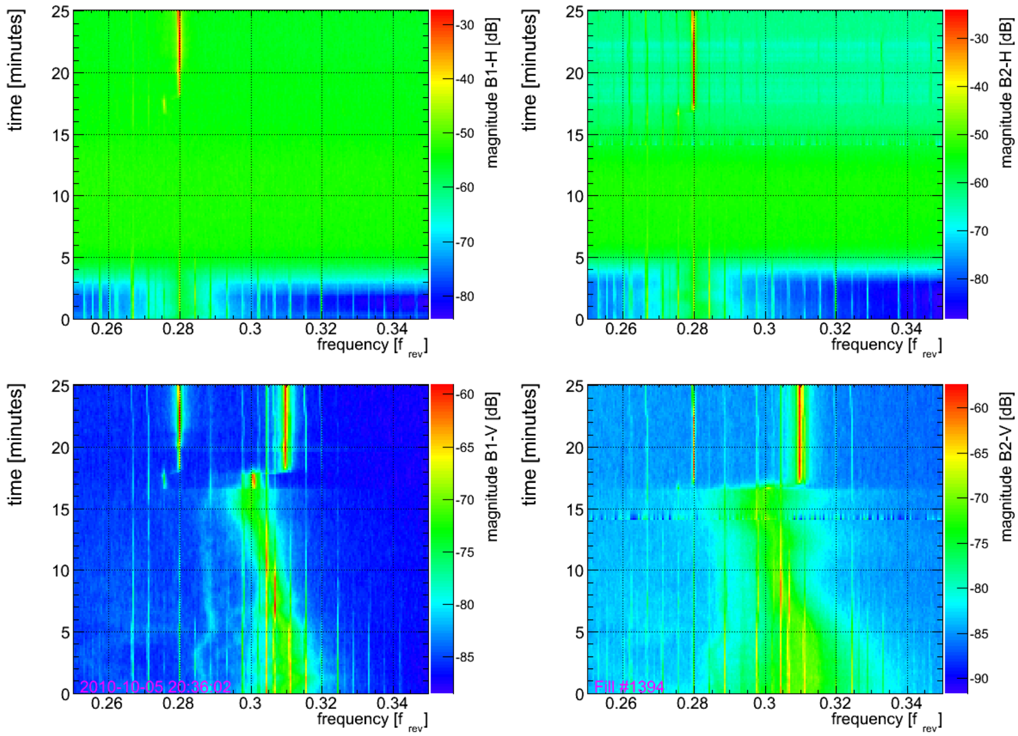

Figure 4: Tune spectra during the ramp of fill 1394. The resonant beam excitation at the higher-order mains harmonics and due to the particular choice of nominal horizontal LHC tune $Q_h = 0.28 * f_{rev} = 3150\,\mathrm{Hz}$ is visible.

loss. Despite the good overall performance and small transmission losses related to Q/Q' and orbit feedbacks, last year's percent-level particle losses may become more critical with the increased stored intensity foreseen, and will continue to be carefully monitored in 2011.

In contrast to initial expectations during the design phase, the higher order beam parameters, chromaticity and coupling turned out to be largely reproducible from fill-to-fill, enforced by strict pre-cycling of the magnets following physics, access and after recovering from errors. The fill-to-fill chromaticity and coupling variations appear to be sufficient controlled for the time being. With the reduced beam intensities reached in 2010 and beginning of 2011, no serious issues have been observed related to beam instabilities. Nevertheless, the conflicting requirements and compatibility between ADT and Q/Q' diagnostics is being further investigated in anticipation of the expected instabilities that may arise to due to electron clouds in the LHC.

ACKNOWLEDGEMENTS

The author expresses his sincere gratitude for the many fruitful discussions, contributions and support during the commissioning given by the LHC operation and commissioning teams and the invaluable support and advise we received while prototyping and testing our feedbacks from our colleagues at BNL, FNAL, DESY, PSI, GSI, Diamond, Soleil, Triumf and many other laboratories. In particular, the contributions of M. Andersen, A. Boccardi, E. Calvo, S. Fartoukh, R. Denz, M. Gasior, W. Höfle, L. Jensen, S. Jackson, R. Jones, Q. King, M. Lamont, E. Metral, S. Page, M. Pereira, B. Salvant, M. Strzelczyk, and J. Wenninger are gratefully acknowledged.

REFERENCES

[1] R. J. Steinhagen et al., BIW'10, 2010

[2] A. Boccardi et al., CERN-BE-2009-002, 2009

[3] R. Steinhagen et al. IPAC'10, CERN-BE-2010-023, 2010

[4] R.Jones,P.Cameron,Y.Luo,BNLCA/AP/204.

[5] Sammut, Bottura et al., CERN-LHC-Project-Report-958

[6] R. Steinhagen, 3rd LHC Project Workshop, 2006

[7] R. Jones, LHC-Performance-Note-006, 2009

[8] E. Calvo, LHC Beam Operation Workshop, Evian, 2010

[9] P. Cameron, BNL, C-A/AP Note 268, 2007

[10] W. Venturini, LHC Beam Operation Workshop, Evian, 2010

[11] M. Zerlauth, LHC Beam Operation Workshop, Evian, 2010

[12] E. Todesco, LHC Beam Operation Workshop, Evian, 2010

[13] "LHC Design Report", CERN-2004-003-V-1, 2004

[14] V. Montabonnet, "Main dipole circuit's mains ripple", private communication

[15] W. Höfle, LHC Beam Operation Workshop, Evian, 2010

[16] G. Arduini, LHC Beam Operation Workshop, Evian, 2010