

FEEDBACKS: STATUS, OPERATIONAL DEPENDENCIES AND OUTLOOK FOR 2011

R. J. Steinhagen, CERN, Geneva, Switzerland


Abstract

This contribution summarises the feedback performance during LHC's first full year in view of higher-intensity operation in 2011. While all involved systems generally performed exceptionally well, this contribution focuses on issues specifically related to operational dependencies and operation of the tune and chromaticity diagnostics instrumentation. Possible mitigation, some of which have been already explored during the year, are being discussed.

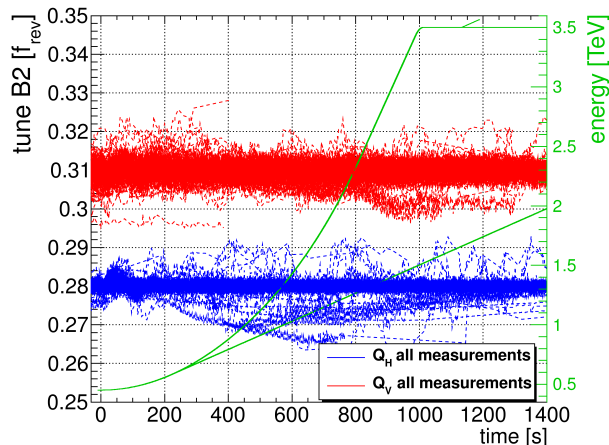
INTRODUCTION

Since the LHC restart in 2010, the Orbit, Q and Q' diagnostics and feedback systems (OFC) were used during almost every fill with the exception of a few ramps used to evaluate the decoupling scheme between chromaticity and the tune feedbacks loop and the few that were affected by outages of the base-band-tune system (BBQ, [1]) discussed below. While the Tune Phase-Locked-Loop (PLL) has been commissioned and used during a few 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 tune and chromaticity from only 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 better than expected (compared to other hadron colliders where similar attempts were made [3]) and soon became the workhorse and base-line mode of operation of the feedbacks.

The feedbacks facilitated a fast and reliable commissioning of the LHC: the orbit-FB kept the largest orbit excursions during the ramp typically below $70 \mu\text{m}$ and down to the residual BPM measurement noise of about $5 - 10 \mu\text{m}$ during the other operational phases. The tune was stabilised typically better than 10^{-3} with initially larger excursions during the snap-back which were further optimised to the same nominal performance. Figure 1 shows the superimposed residual tune stability for beam one and two during 2010. Being used on every ramp 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 due either to direct or indirect feedback involvement, out of which 6 (5) were during the initial 3.5 TeV commissioning.



(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 BBQ outages further discussed in the text.

This contribution summarises some of the feedback issues observed in 2010, the present status of their mitigation measures and possible improvements related to Q/Q' diagnostics and feedbacks in view of 2011 operation.

FEEDBACK ISSUES AND MITIGATION

The few beam dumps related to feedbacks were limited to their initial setup and commissioning during the first months and had a small (below percent-level) impact on overall machine operation [4, 5]. Most of the beam dumps where feedbacks were involved were due to either false-positive QPS trips which have been mitigated by introduc-

ing 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. Some other software error handling of exceptional conditions and common to all feedbacks ('NaN' user-reference and input data, energy transmission errors over the timing system, etc...) were identified timely and fixed by the end of July. Since August, the remaining issues were mainly related to instrumentation quality and integration, such as:

- systematic effects related to the stability of the BPM measurements, discussed in [2],
- interferences of the nominal transverse bunch-by-bunch feedback operation (ADT) with the tune diagnostics, discussed below,
- kernel software updates and denial-of-service security scans of the operational OFC machines during beam operation (causing some beam dumps and down-time) which are necessary but which are to be scheduled during technical stops in the future, and
- integration and automation of reference changes and feedback operations via the operational sequence,

all of which are being addressed in view of the upcoming 2011 operation.

BBQ Diagnostics Outages

Intrinsic to all feedback systems, the ultimate performance of any such system is determined by the performance and reliability of the initial measurements they are based on. In order to reduce the residual dependence of bunch-length and -shape oscillations, a 400 MHz low-pass filter has initially been installed prior to the BBQ to further improve the (in-)dependence of the measured spectrum on longitudinal effects. While this scheme worked well initially for beams with single or a few sparsely distributed bunches, the detector became more sensitive to longitudinal effects with increasing number of bunches. The tune signal-to-noise ratio reduced with every bunch added up to the point (about 50 bunches) where it completely vanished within the noise, subsequently thwarting a reliable tune diagnostics and consequentially feedback operation as illustrated in Figure .

Fast intra-bunch shape measurements performed with LHC's head-tail monitors indicated that the time when these outages occurred coincided with periods of increased longitudinal activity of bunch shape oscillations, a side-effect of the required longitudinal blow-up during the ramp. At the same time it was found that the 'single-bunch peak-detection characteristic' of the BBQ is only valid for bunch

filling patterns beyond about 50 bunches. Below this number, the detector can be sensitive to coupled bunch-to-bunch modes and intensity variations.

In response to this, the BBQ has been reverted to the initial detector scheme, removing the low-pass filter prior to the BBQ detector (since it reduced the effective Tune S/N by about 6 dB but had a minimum impact on the sensitivity on longitudinal effects) and improving the high-voltage rating of some of the components, necessary due to the higher voltage and power-requirements without the low-pass. After this modification, the original sensitivity and spectral performance was restored to some degree, as shown in Figure . The tune signal-to-noise ratio improved by more than 6 dB reducing the impact of the remaining longitudinal activities. This also indicates that an important part of the signal that the BBQ detects, is derived from oscillations that are above 400 MHz (aka. head-tail motion). However, though reverting to the previous scheme helped, the exact mechanism of the original issue (driving source of the head-tail motion, etc.) is still not fully understood and should be closely monitored while increasing the number of bunches and intensities in 2011.

FEEDBACK AND Q/Q' DIAGNOSTICS PERFORMANCE

For the first year of operation, the LHC performance supported by many feedbacks is impressive and transmission losses could be kept below 3% for most ramps. However, these percent-level losses could become more critical for the planned ramp-up of nominal intensities in 2011. A fill-to-fill overview of the evolution of the stored intensities, transmission losses, peak-to-peak tune stability and corresponding required feedback trims is shown in Figure 3. The steady increase in stored intensity per fill is visible. Two markers were added to separate a) the initial commissioning periods of establishing first injection, ramp, squeeze and collisions with low intensity beams, b) operation with nominal proton bunches and later bunch trains, and c) ion operation. Most losses occurred when switching mode of operation e.g. changing from single bunch injection, to trains and to ion operation. Some (scraped) halo losses have been seen but it is believed that these particles would have eventually been lost in the collimators anyway, and for the few ramps and periods where radial modulation were applied systematically to measure $Q'(t)$, little or no direct impact of the modulation ($\Delta p/p < 2 \cdot 10^{-4}$) could be seen on transmission losses or beam size growth.

Tune-FB Stability

The tune-FB performance was fairly steady over the fills and largely dominated by the snap-back as shown in Figure 1. A direct decay of main quadrupole currents or feed-down effects coming from the main dipole's b3 decay could be the cause of these variations, as discussed in [6].

Initially, very conservative feedback settings were cho-

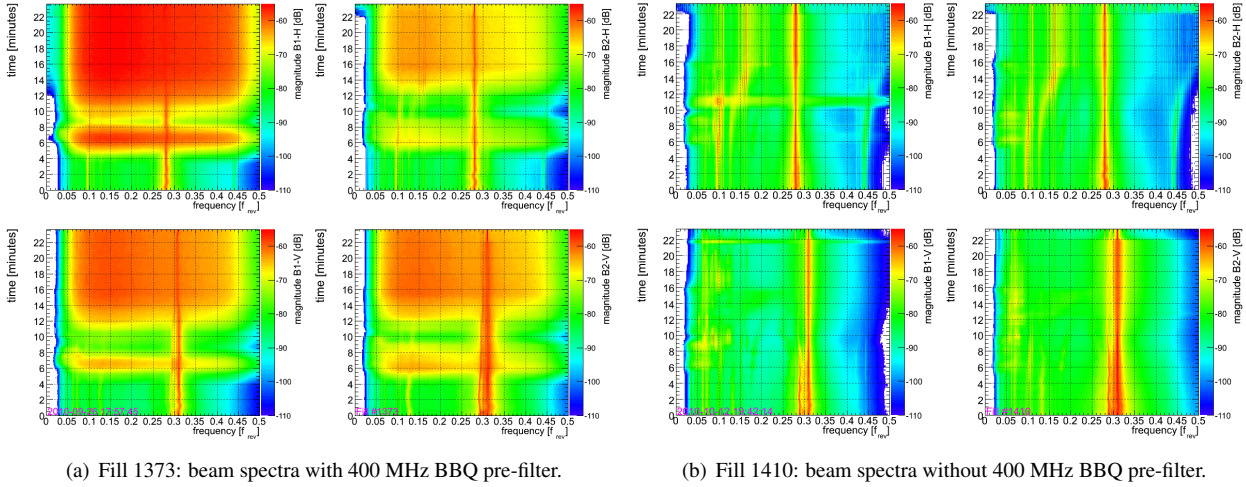


Figure 2: BBQ outage during ramp before and after the 400 MHz low-pass filter removal.

sen, which resulted in exceeding the initial tune stability requirement 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 $3 \cdot 10^{-3}$ as visible in Figures 3(e) and 3(f). In any case, the stability is limited by the resolution, stability and reliability of the Q/Q' diagnostics rather than the feedback controller or loop itself.

Operational Dependence on Feedbacks

As visible in Figures 3(g) and 3(h), the corresponding tune trims rather increased than decreased over time which correlated with the progressively reduced frequency with which the systematic dynamic real-time tune trims were incorporated into LSA's static feed-forward function. Also, for some fills the real-time trim action substantially exceeded the typical correction range compared to previous fills. In these cases, the feedback compensated for effects that were introduced either directly (human and/or incorporation errors) or indirectly through feed-down effects that were otherwise not accounted for by the day-to-day operation (such as incomplete pre-cycles after accesses, newly measured $Q'(t)$ incorporation into the ramp functions). These examples nicely demonstrate that – even with perfect feed-forward incorporation of the recurring real-time actions – feedbacks can and did provide some additional safety margin to operation by indifferently suppressing and absorbing unexpected perturbations. At the same time, it should be pointed out that the beams without feedback support would have been probably lost 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 [9]. In order to reduce this dependence on feedbacks, which is the mandatory requirement to have them fully operational and always operating at with nominal performance for every fill, it is strongly recommended to systematically monitor and transfer recurring real-time feedback actions into the ramp and squeeze functions.

Chromaticity Stability

While the availability of the intensity, tune and feedback trim data is extensive and generally available for nearly every fill, the data on beam size evolution and in particular $Q'(t)$ is very sparse. However, for the few consecutive fills for which $Q'(t)$ was measured indicated a fairly reproducible behaviour as shown in Figure 4. A first order magnetic field correction of the chromaticity has been applied to the ramp using the MCS spool pieces. The remaining largest fill-to-fill variations occurred as expected during the first 200 seconds of the ramp reaching up to $\Delta Q' \approx \pm 5$. Once reaching 3.5 TeV another decay of about 6 units of chromaticity is evolving and to allow this decay to settle, the ramp was artificially extended by about 6 minutes. In between the chromaticity was found to be stable within about $\Delta Q' \approx \pm 2$ which indicates that beside the snap-back most of these effects could be compensated by a feed-forward function nearly down to nominal requirements. Still, all ramps exceeded the initially required chromaticity stability of $\Delta Q' = 2 \pm 1$, often with systematically negative chromaticity as can be seen in Figure 3(j). While the effect of operating with negative chromaticity was partially absorbed by the ADT, it is recommended – similar to the tune perturbations and feedback – to correct for this systematic effects to reduce the unnecessary systematic dependence on feedbacks.

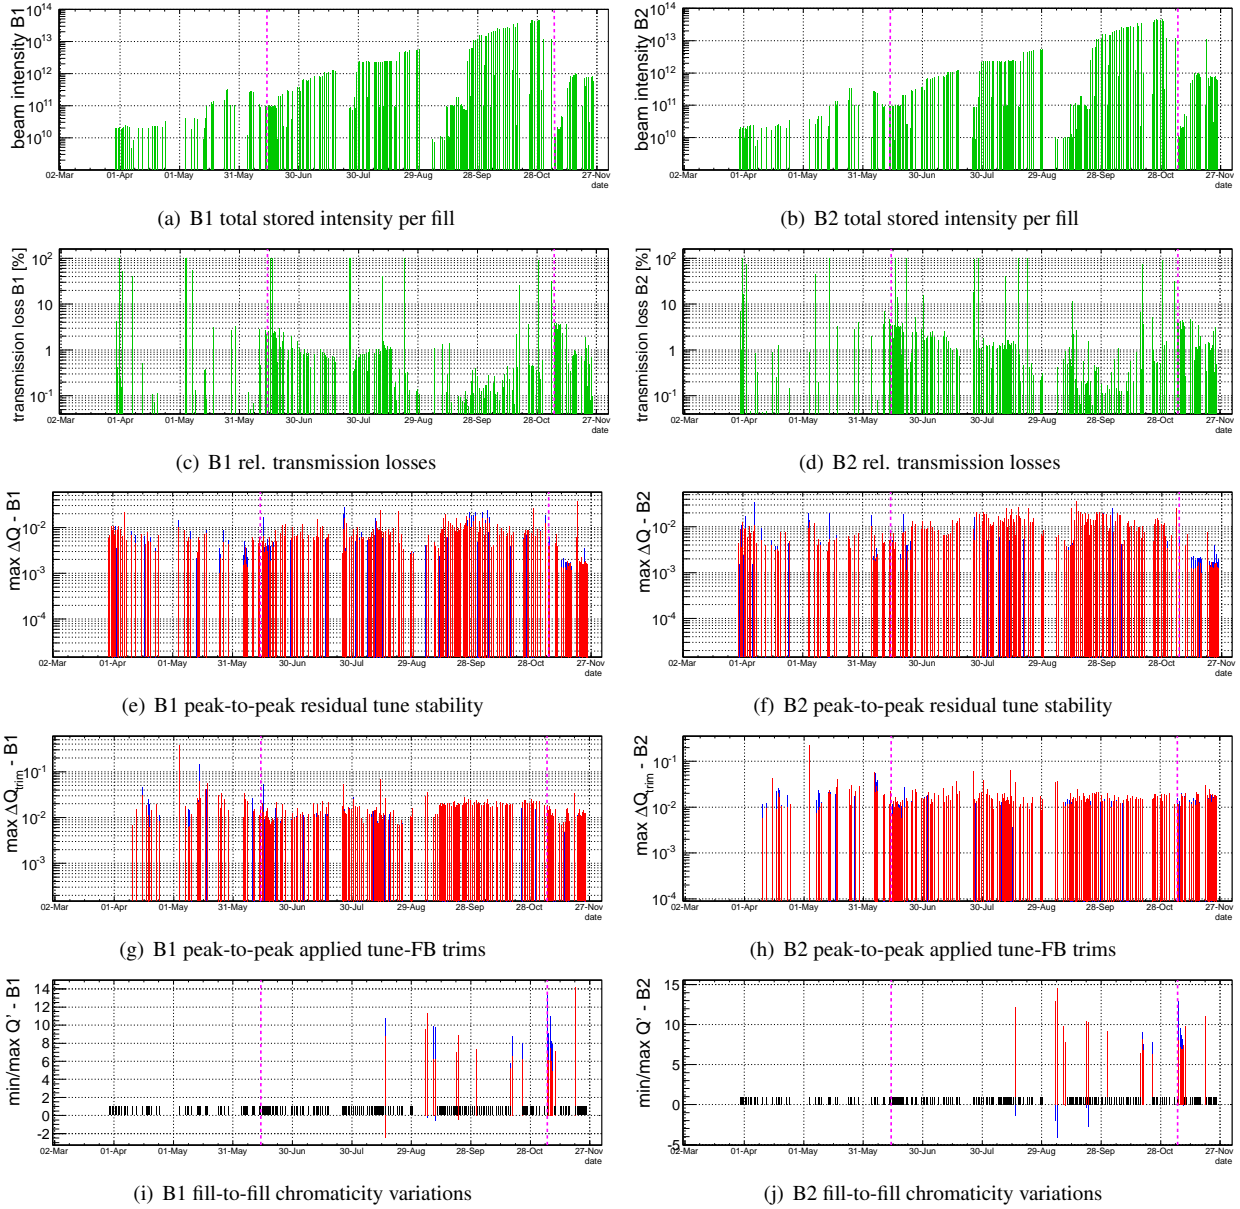


Figure 3: Q/Q' -related fill-to-fill performance overview of 2010: evolution of the stored intensities, transmission losses, peak-to-peak tune stability and corresponding required feedback trims. The two magenta markers indicate the two major changes of mode of operation: a) from initial commissioning to gradual intensity increase and b) from proton to ion operation. Parameters related to the horizontal plane are indicated in blue and for the vertical plane in red.

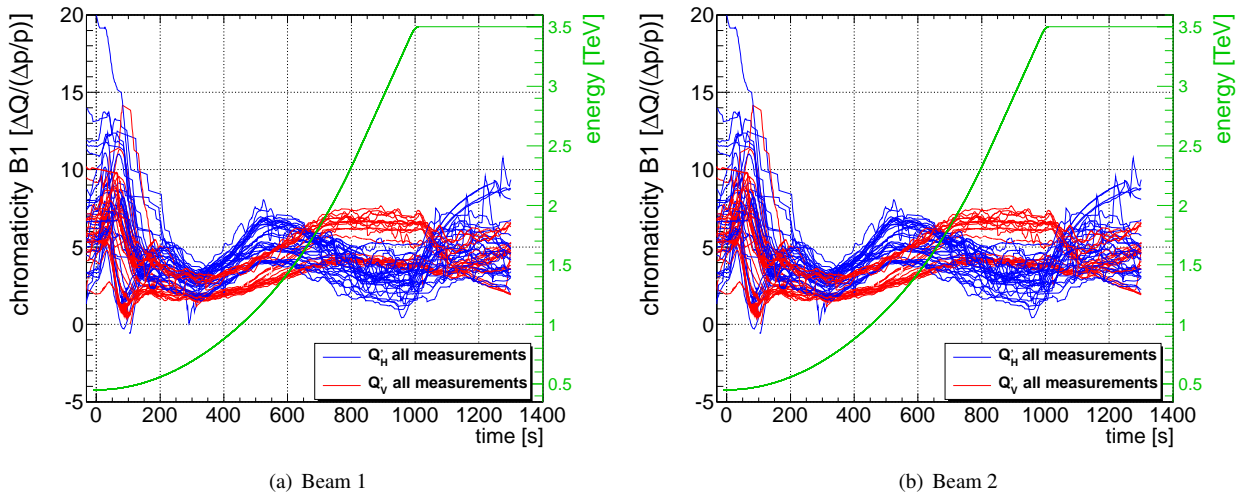


Figure 4: Residual superimposed $Q'(t)$ stability during the ramps in the time periods indicated in Figures 3(i) and 3(j).

Impact of Q/Q' Stability on LHC Operation

The actually observed Q/Q' perturbations are in good agreement with the expected perturbations and initial design assumptions [7]. With exception of the measurement and control of $Q'(t)$ that – while the diagnostic and feedback was available – has been given less priority, all parameters could be kept just above the initially targeted limits. To be further investigated: are this slight out-of-tolerance parameter stabilities acceptable for operation with nominal beam, or equivalently, is the achieved feedback performance adequate? Or does it require further improvement?

Thus an extensive analysis of transmission losses and beam size growth as stability indicators with biggest impact on luminosity production has been performed, to assess the impact of feedback performance on operation. Since the largest and fastest tune and chromaticity variations occur during the ramp, the presented analysis focuses on a total of 275 ramps for which the given parameter, feedback actions as well as the beam stability indicators were available. In this analysis, the transmission loss is defined as intensity loss between the start and the end of the ramp, excluding the loss of un-captured beam at the very beginning. As discussed in [8], for the analysed period, neither the synchrotron light nor ionisation profile monitor could provide reliable beam size growth measurement during the ramp. In order to nevertheless assess some form of fill-to-fill beam size growth changes, the beam sizes at injection were compared with those at flat-top, including some best effort correction factors which were constant over the analysed periods. While this does not provide an absolute measure of the relative beam size growth, it remains a rough indication whether the beam size changed during the ramp from a fill-to-fill perspective. The corresponding correlation plots are shown in Figure 5. Comparing the individual stabilities during the ramp on a fill-to-fill basis seem to indicate an (anti-)correlation between 0.5 and 0.7 between the residual peak-to-peak chromaticity variations and transmission

loss and beam size growth. Thus, the higher the chromaticity swing during the ramp, the less particles are lost but also the larger the beam size growth. This result would to first-order relate well with expectations of the beneficial effects of large(r) chromaticities on collective instabilities and detrimental with respect to higher-order head-tail modes causing emittance blow-up. While the statistics supports the case of Q' -related transmission losses, the effect on beam size growth, in particular the absolute magnitude, remains substantially limited by the systematic errors on the beam size measurement. In order to assess the full magnitude of this effect, it would be useful to further explore this effect through a controlled experiment at constant energy for which both the synchrotron light monitor and ionisation profile monitor provide better beam size estimates.

MAINS HARMONICS

As can be seen in Figure 5, no direct correlation between residual tune stability and beam size growth is visible. There is a limited correlation between residual tune stability and intensity transmission during the ramp, with the 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 for example in Figure 6. A set of mains harmonic are 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[10]. Their impact is a priori not a big issue and similar to evading the 'hump' could easily mitigated e.g. by shifting the nominal working points by 0.001 only.

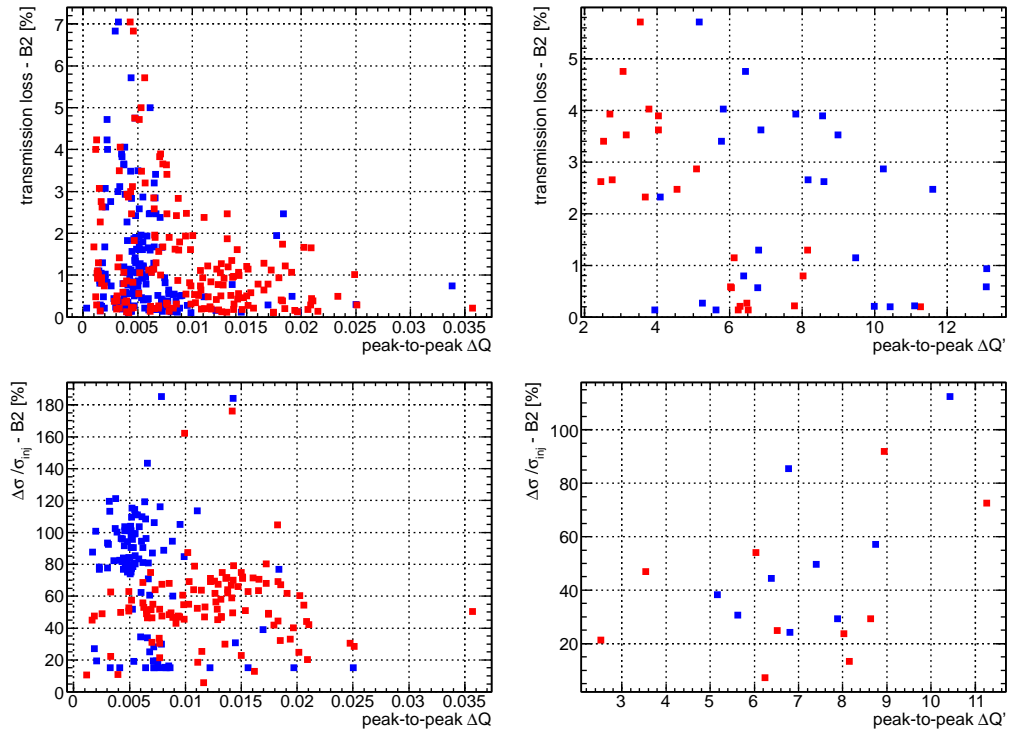


Figure 5: Correlation plots showing transmission loss and relative beam size growth versus peak-to-peak tune and chromaticity stability during the ramp. Beam 1 (blue) and Beam 2 (red) are indicated. The relative beam-size growth should be interpreted only as a linear measure of the fill-to-fill variation. At the time of the analysis there were still significant uncertainties on the synchrotron-light and BGI based beam size measurements with strong uncertainties on the absolute scale – however the scale being reproducible from one fill to another fill.




Figure 6: 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 $Q_h = 0.28 * f_{rev} = 3150$ Hz is visible.

COHABITATION OF ADT AND Q/Q' DIAGNOSTICS SYSTEMS

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 in [9].

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[1]. 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 [11], BPM: 50 μm [2]). 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-metre level. 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 originating from mains and ADT, the 'hump', and other effects documented elsewhere[11, 12].

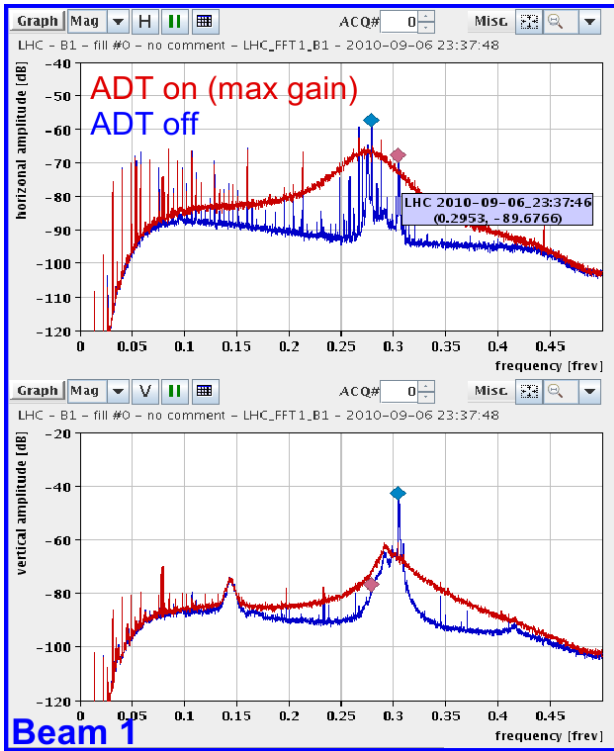
While these μm -level oscillations are a-priori beneficial for a passive detection of the tune, they are incoherent 'noise' from a FFT or PLL diagnostic point of view. Regardless of whether using a driven FFT- or PLL-based diagnostic tune system, the beam needs to be excited about 20-30 dB above this 'noise' to recover the same reliable performance as using residual oscillations only. The corresponding absolute amplitude of about 10 – 100 μm that is excited on top of the residual tune oscillations are in conflict with collimator requirements (< 200 μm 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

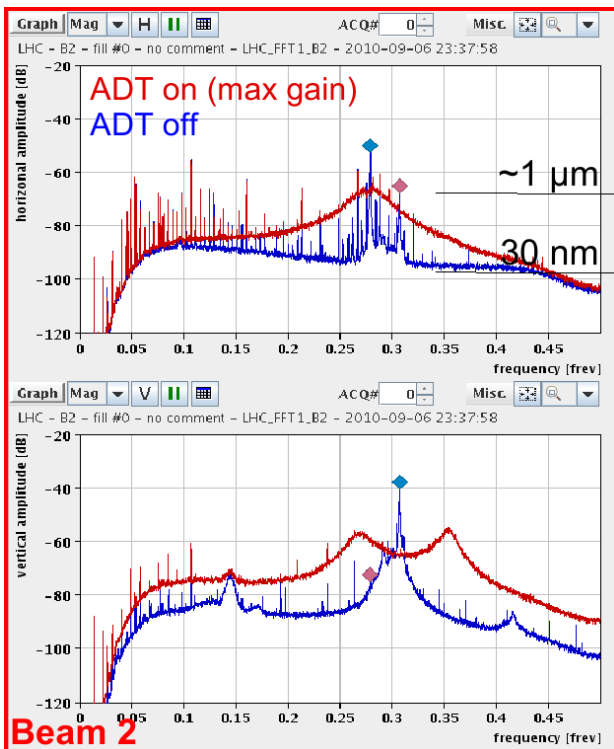
The ADT has been successfully operated since July, damping injection oscillations on a regular basis, and being kept 'on' also during ramp and during collisions with an impressive performance of damping times of few hundred down to 50 turns[11]. At the same time and, as one of the limiting factors of any feedback, part of the ADT measurement noise is propagated onto the beam as illustrated in Figure 7, compromising the BBQ high-sensitivity capabilities by up to 30 dB and reducing the tune resolution by at least two orders of magnitude. By comparison of the unperturbed and damped spectra, the particular shape of the noise probably originates from the particular internal ADT filters and feedback gains, and in many cases, the maxima being unrelated to the actual tune-resonance. In addition, the ADT – used as an abort gap cleaner – creates ringing excitation. This ringing prevails up to 250 ns and resonantly excites e.g. the first bunch after the abort gap with the given frequency that does not necessarily correspond to the tune. This effective ADT-induced noise floor and observed bunch-to-bunch cross-talk hinders, and in some cases, prevents reliable operation of LHC's Q/Q'-diagnostics and related feedbacks.

Some mitigation options – of which some have been tested in 2010 – that could make the Q/Q'-diagnostic compatible with the ADT function are:

1. low(er) ADT gain after injection until end-of-squeeze
2. high-ADT gain for first N-turns after injection, then lower-gain
3. sacrificial (e.g. non-colliding) bunch for which ADT is disabled or operated at a low-gain
4. dead-band in ADT gain function which masks oscillations slightly above its BPM noise floor
5. deriving tune from ADT's residual exciter signal
6. operating with high ADT gain and Q-PLL exciting about 30 dB above ADT's noise floor. This option was tested during 2010 but was found to be impractical because of the measurable emittance blow-up, particle loss and complex dependence on ADT gain, energy, intensity and other collective effects.
7. operating with high ADT gain and Q-PLL exciting about 30 dB above ADT's noise floor. This option is similar to the previous one, but preferred since the excitation levels are less critical and on the 10 μm . However the technological feasibility of this noise reduction needs to be demonstrated.
8. operating with high ADT gain and deriving the Q/Q' signals from the transverse Schottky monitor, methods involving off-resonance and/or exciting outside the ADT bandwidth



(a) Beam 1 spectra



(b) Beam 2 spectra

Figure 7: Comparisons 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.

The first two options are presently the only viable, reliable and available options until the end of 2010-11, the second differing just by the ADT being adapted to changing requirements. The third option cannot be exploited for the time being due to the afore-mentioned ADT ringing and lack of bunch selector capabilities for the BBQ. The latter would require further research and development to not compromise the existing system's signal-to-noise performance and reliability. Beside the first two options, all have in common that besides some additional simulations and hardware development, all are 'long shots' and require more operational and long(er)-term experience with respect to robustness, resolution and bandwidth prior to being used within the Q/Q' -feedbacks.

PLANNED FEEDBACK MODIFICATIONS

Most of the modifications planned for 2011 are minor, limited to communication protocols, additional logging requirements and clean-up of dead-code or functionalities that have been implemented but found to be unused or unnecessary during day-to-day operation.

The handling of dynamic orbit reference has been in place since 2008 but needs further testing and integration into LSA, sequencer and operational GUIs (YASP). This integration should, for the time being, also eliminate the frequently used but error-prone masking of BPMs during squeeze that 'blind' the feedback with respect to dynamic changes inside the insertions. The disabling was an effective workaround, but providing the OFC with shape and time-evolution of the changing reference is the cleaner and more reliable solution.

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 otherwise reduces the noise that is propagated from the beam instrumentation to the beam via the feedbacks (e.g. during collisions). 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 control the 'high' and 'low' extremes of bandwidth.

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. Urgent issues have been resolved in a timely manner, and (less critical) systematic BPM and Q/Q' performance issues are being followed up in view of nominal LHC operation. Analysis of the feedback actions of more than 280 logged ramps indicated that more than half of all fills would have been lost without feedback support and the others likely affected by

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

The measurement and control of $Q'(t)$ received less attention than was initially planned, with systematic negative chromaticities and large relative variations during the ramp. The few measurements performed during the ramps indicated an intrinsic trade-off between beam stability (and low transmission losses) and beam size growth as a function of chromaticity. There are still some important uncertainties on the absolute scale of this effect and it would thus be useful if these dependencies could be assessed in more detail during controlled measurements at injection and top energy.

If operated at maximum gain, the effective ADT-induced noise floor and observed bunch-to-bunch cross-talk of the current abort gap cleaner implementation hinders reliable operation of LHC's Q/Q' -diagnostics and related feedbacks. Mitigation options compatible with a high-gain ADT operation will be further explored in 2011. At the same time, the indifferent high-gain ADT operation should be validated against the actual instability growth times to optimise the required damping constants against the noise that is propagated on to the beam.

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. In particular, the contributions of M. Andersen, A. Boccardi, S. Fartoukh, M. Gasior, W. Höfle, S. Jackson, R. Jones, M. Lamont, E. Metral, B. Salvant and J. Wenninger are gratefully acknowledged.

REFERENCES

- [1] A. Boccardi et al., "An overview of the LHC Transverse Diagnostics Systems", CERN-BE-2009-002
- [2] E. Calvo, "BPM: status, measurement reliability and outlook for 2011", these proceedings
- [3] P. Camperon, "RHIC run 7 Feedback Ramp Tunes", BNL, C-A/AP Note 268, 2007
- [4] W. Venturini, "Operational efficiency", these proceedings
- [5] M. Zerlauth, "Do we understand everything about MP system response?", these proceedings
- [6] E. Todesco, "Can we improve the magnetic model/cycle and its effects?", these proceedings
- [7] R. Steinhagen, "Real-Time Feed-Forward/Feedback Required", 3rd LHC Project Workshop, 2006
- [8] F. Roncarollo, "Can we get a reliable on-line measurements of the transverse beam size?", these proceedings
- [9] "LHC Design Report", CERN-2004-003-V-1, 2004
- [10] V. Montabonnet, "Mains ripple on the LHC main dipole circuits", private communication
- [11] W. Höfle, "Transverse damper operation", these proceedings
- [12] G. Arduini, "Hump: how did it impact the luminosity performance", these proceeding