EXPERIENCES WITH FEEDBACK SYSTEMS AND FORESEEN IMPROVEMENTS FOR LS1

Ralph J. Steinhagen, CERN, Geneva, Switzerland

Abstract

This contribution summarises the impact of the LHC real-time feedback systems on machine protection and availability. The effects leading to combined failure modes causing beam losses, and those stressing the machine protection system, as well as their planned and proposed mitigations during and after LS1 will be discussed.

INTRODUCTION

The Large Hadron Collider (LHC) at CERN deploys a comprehensive suite of beam-based feedbacks (FB), which not only improves the performance of machine operation but also ensures the proper functioning of its machine protection and beam cleaning systems [1, 2, 3, 4, 5, 6, 7]. Subsequent improvements based on the gained experience with LHC beam are described in [8, 9].

Feedbacks on their own do not directly impact on LHC machine protection. However, incorrect or the absence of FB corrections may contribute to dangerous failure scenarios that, combined with another fast failure incident, may compromise machine protection. Two of the more frequently discussed possible failure scenarios are: a) local orbit bumps combined with, for example, a fast kicker failure, which may violate the global collimator hierarchy, and subsequently cause severe single-turn losses; b) insufficiently controlled tune or chromaticity deviations that, together with another beam instability incident, could unfavourably change the time-scale and distribution of particle losses inside the machine.

LHC FEEDBACK ARCHITECTURE

In order to cope with the large numerical complexity, 'firm real-time' (RT) requirements, the reliability and the availability aspects of high-performance computing hardware, the central feedback controller has been split into two functional parts as illustrated in Figure 1: the Orbit

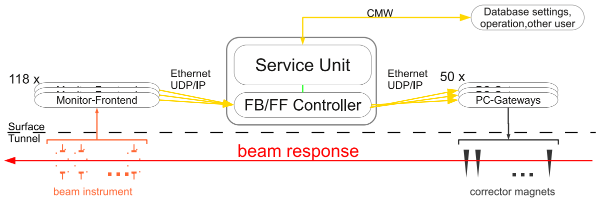

Figure 1: LHC Feedback Schematic.

Feedback Controller (OFC), responsible for the nominal FB controller functionality; and the Orbit Feedback Service Unit (OFSU), responsible for the higher-level feedback functionalities. These higher-level functionalities are setting management, beam parameter reference handling, triggering of feedback actions in response to external timing events, and general integration into the operational environment, e.g. relaying data to clients, interfacing to LHC experiments, logging infrastructure, triggering the software machine interlock system, control room interfaces, etc.

While the OFC and OFSU are the most visible single feedback sub-systems, the whole loop consists of more than 3500 devices (controlled via about 130 front-end computers), and depends on many technical infrastructure service i.e. the FESA front-end framework, CMW communication middleware, the timing system, and the technical network. Thus, the strength of the overall loop depends on the reliability of its weakest link – not necessarily always being the OFC or OFSU servers.

In order to achieve a maximum robustness and closed-loop performance, the OFC has been designed as a simple input-processing-output streaming server, utilising only a very limited number of concurrent threads. The reason for this restrained use of threads was to reduce the variance in dynamic load in favour of a constant periodic load, which is by-far the easiest RT paradigm to quantify and guarantee the compliance with the given real-time constraints. Consequently, conditional statements and semi-random/user-driven code execution are delegated to the OFSU in favour of keeping a predictable constant load on the OFC. More generally, most of the available software technologies (notably Java), algorithms or libraries that could not be thoroughly qualified in terms of worst-case RT latencies were omitted from the OFC.

EXISTING FB ERROR MITIGATION

FBs may amplify the effects of other failures, however, they mostly impact machine availability for physics. Most of the deteriorating effects linked to FB operation that may lead to combined failures are typically slow. They are detected or mitigated by either:

A) the feedback controller itself, with verification time-scales in the order of typically 40 to 80 ms. The large majority ($\approx 80\%$) of the OFC functionality thus deals with error handling rather than the main feedback logic ($\approx 20\%$): receiving the beam instrumentation data, computing the corrections that minimise the deviation of the given beam parameter, and sending the updated reference settings to the power converters and RF function generators (FGCs). In order to achieve an equivalent analogue bandwidth of up to about 1 Hz, the feedback controller handles the input and output data from about 118 front-end computers via a Gigabit-Ethernet backbone using UDP/IP at a sampling rate of typically 25 Hz. Most of the OFC error handling is related to sanity checks of the data integrity and the assessment of the beam parameter stability. This information is subsequently used to automatically de-select (or momentarily mute) an erroneous beam instrumentation input or – if necessary – to stop the FB loop. Many of the involved thresholds are subject to tuning, in particular the trade-off between availability and spurious quench protection system (QPS) trips that may occur if the FB acts on noisy input data.

B) the Software Interlock System (SIS), with surveillance and reaction time-scales in the order of seconds. The SIS monitors and provides slow interlocks on the global orbit, RF frequency, and closed-orbit-dipole currents. It also monitors the OFC/OFSU status to catch erroneously latched references. It also preventively dumps the beams in case of a loss of CMW communication with the front-ends (watch-dog time-scale: 10 seconds). Through the SIS, since many of these software interlocks depend on the OFC and OFSU functioning, the overall FB loop reliability and availability effectively becomes an interlock.

For the time being, the tune and chromaticity measurement values are not interlocked. However, most of the related failure scenarios are indirectly interlocked through the monitoring of the losses seen by the Beam Loss Monitors (BLM). The information related to issues regarding UDP latencies (missing packets, bursts, etc.), CMW communication, technical network latencies and timing infrastructure is already being monitored by the OFC but not yet further exploited (e.g. as forewarning). While some of these symptoms are not necessarily failures, they are often indicative of non-ideal situations that – if aggravated – could lead to a front-end or system failure and subsequent dump of the beam.

FB FAILURE STATISTIC

Overall, the FB performance was good. This proved to be essential for nominal LHC beams, in particular the Tune-FB during the ramp start (snap-back) and Orbit-FB during the beta* squeeze. The large majority of fills were made with minimal losses into physics [10, 11, 12]. Nevertheless, a few fills were lost due to non-nominal FB behaviour, which deserves careful scrutiny to identify and possibly improve systematic or recurring errors. As discussed in [13], based on the analysis of 131 beam dump post-mortem events in 2011 occurring during the operational 'ramp' and 'squeeze' phases where the feedbacks are nominally 'on', 33 dumps were attributed to feedback systems (Orbit- and Tune-FB combined). Most of the dumps (23 out of 33) were related to QPS triggers in response to noisy Tune-FB real-time trims (notably for the RQTH and RQTF tune trim magnet circuits). As a mitigation, the very tight QPS thresholds were raised for 2012 from 0.1 to 2 V and the noise performance of the beam instrumentation (BBQ) feeding into the Tune-FB has been improved. Based on these mitigations, in 2011 it was estimated that two to three FB-induced beam dumps were encountered during the LHC operation in 2012/13.

A similar post-mortem analysis has been repeated for the 2012/13 operational period. Table 1 shows the summary of the analysis. The numbers take into account only post-mortem events with energies above 450 GeV and beam intensities larger than 10^{10} protons per beam to exclude special MD-type experiments with pilot beams (PM comment key-words selection criteria: 'FB', 'Feedback', 'OFC', 'OFSU', 'BBQ', 'BPM', 'RT', 'Orbit', 'Tune', 'Instability'). The statistic includes only dumps, i.e. near-misses, events causing losses without a dump, or events that have been recovered by the operations crew or the sequencer are excluded. As can be seen in Table 1, the overall number of post-mortem events increased.

Table 1: Summary of FB-induced beam-dumps during 2010-2013 LHC operation.

	Total PM	FB-induced	Percentage
2010	453	8	1.7%
2011	684	30	4.4%
2012/13	851	28	3.3%

To first order, the number of FB-induced dumps did not decrease as predicted in [13] but was steady, indicating another new set of failures. Table 2 gives an approximate break-down of the FB-induced post-mortem events according to the FB sub-systems. Please note, some post-mortems are counted twice as the underlying effect could not always be unambiguously disentangled between the various sub-systems. The numbers are thus indicating trends rather than being absolute. As predicted in [13], it can be seen that the vast majority of dumps related to the interplay between the QPS and BBQ signal integrity reduced significantly during 2012/13. However, a new source of equipment failure emerged related to problems with infrastructure, over which the equipment owners have only limited control (i.e. FESA, CMW, timing, technical network).

The main causes leading to FB induced beam losses or to stress of the machine protection system, can be grouped into three sub-categories:

- measurement quality (BPMs, BBQ), causing transients on orbit or tune that subsequently provokes losses through pushing the beam on the collimator or inducing triggers of the QPS.
- front-end or software infrastructure problems (OFC, OFSU, FESA, CMW, Timing & network). Many of these actually related to threading issues exhibiting non-RT behaviour, front-end crashes due to memory leakages or out-of-bound accesses, and external

Table 2: Break-down of FB-induced beam-dumps during 2010-2013 LHC operation according to FB sub-systems. The last two columns indicate PM events more related to orbit and tune related instabilities. The latter are not necessarily related to feedback operation but illustrate the increasing criticality of orbit and tune control.

	FB induced total	OFC	OFSU	BBQ	BPM	QPS	Orbit-Instability	Tune-Instability
2010	8	2	0	2	0	3	9	0
2011	30	2	5	18	3	14	13	6
2012/13	28	4	10	1	7	1	17	30

load factors such as slow clients blocking the BBQ or OFSU communication on front-end machines, and technical network switch overloads that suppressed operationally critical data streams. For example the non-RT behaviour of the input data stream, in particular for the BPM front-ends, has a fundamental impact on the feedback function as the absence of data causes the FB to accumulate latencies. These may either cause the beam parameter to drift and potentially push the orbit towards the collimators, or in some cases cause 'classical self-amplifying FB instabilities' when the correction is applied out-of-phase.

- insufficient loop stability margin, with the FB running at or beyond the design stability limit due to e.g. a mismatch between actual optics and the one used by the OFC.

Compared to 2011, some of the failures became more notable due to the newly added SIS interlocks in 2012/13 and the significantly tighter beam stability requirements and stability margin, as can be seen also in the last two columns of Table 2, which indicates the number of orbit- and tune-instability driven beam dumps.

FORESEEN FB IMPROVEMENTS

Many of the recurring or systematic errors (i.e. OFSU memory leaks and out-of-bound reads) have already been addressed during the 2012/13 operation and most of the dumps occurred before September 2012. However, there remains a number of additional improvements for after LS1, that couldn't be deployed during 2012/13 beam operation.

Measurement and Data Integrity

- Temperature stabilised BPM racks, which should minimise most of the temperature related beam position measurement drifts. If necessary, for a few dedicated pick-ups the remaining drifts could be further allayed by additional RF commutation switches at the pick-ups that can help to identify and compensate measurement errors w.r.t. real orbit drifts, as e.g. already installed for the BPMSW pick-ups closest to IP5.

- During LS1 and LS2, it is planned to deploy redundant read-out electronics for all IR-BPMs based on the positive experience with the newly developed high-accuracy Diode-Orbit acquisition system [14, 15]. The initial deployment will be done only for a selected number of BPMs close to the experimental IPs (BPMSW.1[L/R][1,5,8,2].B[1/2]) rather than the full deployment for all BPMs between cells Q1 to Q7. The naming convention of the additional channels needs to be addressed, as well as their integration with respect to the existing standard LHC BPM acquisition electronics.

- The yet to be deployed new TCTP collimators with integrated BPM buttons will also use the Diode-Orbit acquisition system [14, 15]. However, their integration w.r.t. the regular BPM data, reference management and possible future use within the FB loops needs to be evaluated.

- The Tune-FB presently uses an algorithm that only tracks a single, highest peak with a given minimum line width and within a preset tune frequency range. However, the BBQ spectra often shows multiple peaks in the tune range for nominal LHC beams, which sometimes hampers the correct tune detection as the desired tune line is not always the most dominant peak. In order to improve and avoid locking on spurious, non-tune-related peaks in the BBQ spectra, the possibility of tracking multiple peak candidates is being investigated. The relationship between multiple peak candidates could be used to distinguish between interferences, synchrotron side-bands, and the correct tune eigenmodes.

- In order to complement the existing BBQ based tune diagnostics, it was proposed to integrated the ADT transverse bunch-by-bunch feedback system as an alternate source for tune and chromaticity feedbacks [16].

Improvement of Loop Stability

- The operation and performance of the LHC FBs depends not only on the validity of the amplitude of the magnet current changes but equally on the time and latencies at which they are applied. The system can tolerate occasional latencies but reacts adversely if the 'firm real-time' constraints on the loop delay are not met for multiple subsequent samples or for a few seconds. Investigations indicated that under certain (unknown) conditions, the arrival time of the BPM information does not always comply with these requirements, causing a loss of loop stability and an increase in beam losses. The BPM/BBQ UDP transmission robustness and their implementation needs to be reviewed during LS1, in particular the interplay between CMW, FESA, CMW proxies and technical network infrastructure.

- In order to mitigate possible data congestion on the technical network it is planned to decouple the RT traffic used for the FB from those needed for operation and other services using the existing quality-of-service functionality of the installed network routers and switches. The robustness and reliability of this possible solution needs to be verified.

- For most of 2010-2013, the OFC operated with a single and only later with two beam optics configurations (the second being used for the squeeze). This mode of operation was acceptable for the initially very low FB bandwidths but became increasingly inefficient and reduced the stability margin for the higher bandwidths required during the later 2012/13 operation. Thus, for post-LS1 operation it is favourable that the OFC tracks and uses the actual optics of the machine more closely, particular during the squeeze.

- While the basic infrastructure for gain scheduling is already available, it is not yet used during regular operation. Higher bandwidths are typically only required during a very specific and brief operational period (e.g. start of the ramp, or during the beta-star squeeze) and could be reduced in favour of a more robust loop behaviour. This feature remains to be integrated into the operational environment.

- Most of the tracking and validation of the BPM functionality was done manually and thus was fairly infrequent in 2012/13. An automated procedure similarly to the BLM procedure that is executed before each fill, would help to improve the system performance. The proposed idea is to perturb the orbit with a given pattern at the beginning of each fill and to record the measured BPM response, noise, etc. The measurement itself requires less than a minute if automated, and – if done regularly – small deviation or drifts from the reference could indicate at an earlier stage that are on the verge of failing.

- Even though some of the orbit perturbations (e.g. during the squeeze) are fairly large and fast, and approaching the design limit of what the FBs can handle, most of the feedback actions are fairly reproducible from one fill to the next. A feed-forward of the recurring corrections averaged over several fills could reduce the effective fill-to-fill orbit deviations and their speed. This may allow to reduce the necessity and criticality of running the FBs at high bandwidths with reduced stability margin. An initial implementation is being prepared [17].

- Many of the specific feedback actions and stability depend on a wide range of external conditions, some as simple as the beam presence, availability and quality of beam input data, and response to orbit dipole corrector current settings. In order to allow new FB schemes and controls integration to be thoroughly tested under safe conditions and also during non-beam operation, the option of having a dedicated full FB test-bed should be revisited. This functionality is of particular importance for training and development during the approaching long shutdowns.

Diagnostics and Tracking

Most major problems have been quickly identified and addressed during the first years of operation. Further improvements would benefit from better and more specific monitoring, reporting and tracking of the underlying technical infrastructure, and monitoring of the FB and its subsystem errors. This would help isolating the original cause of the problem, rather than analysing the collateral symptoms that caused the beam dump. For example:

- a finer granularity of post-mortem reports attributing the errors according to the feedback sub-categories would be useful for the system experts to follow-up a given problem (e.g. FB function, FESA, CMW/communication, timing, network, BPM/BBQ input, etc.).

- a systematic monitoring of the infrastructure (FESA, CMW, timing, network availability/latencies, front-end statuses, etc.).

- both the OFC and OFSU collect and provide a lot of information, but only a fraction is being exposed to the control room.

Most of the required functionalities and tools are already available on an expert-level, but are still largely inaccessible or unintelligible for untrained people and day-to-day use in the control room. A better GUI-level integration would be desirable to communicate conditions that are not yet critical for beam operation but could lead to some down-time or beam dump if ignored or further aggravated.

FB REVIEW CONCLUSION AND RECOMMENDATION

In addition to this MPP workshop, another subsequent dedicated review has been organised at CERN to analyse the architecture of the existing FB systems in view of their consolidation and adaptation for post-LS1 LHC operation [18]. The conclusion and recommendations of the review board[1] as given in [18] are:

- the OFC appears to be essentially robust and good for its designed purpose. However, additional functionality is expected to be required, and a number of specific

[1] M. Lamont (BE-OP, chair), Javier Serrano (BE-CO), Jakub Wozniak (BE-CO), Stephane Degahaye (BE-CO), Quentin King (TE-EPC)

issues identified above and also detailed in [18] should be addressed, notably the establishing of a full test environment before proceeding with any significant new functionality.

- the OFSU is unmaintainable in its present state and a factorisation is required, notably the split between data re-distribution, settings controls and configuration flow.
- the staffing of the system has been identified as a serious issue and must be addressed urgently.

A complete list of specific items, new functionalities and actions are given in [18].

CONCLUSION

Feedbacks have a priori no 'direct' link to the LHC machine protection system but can create dangerous combined failure scenarios. The OFC/OFSU are the prominent components to the overall feedback loop, which itself depends on many more devices and technical infrastructure services. The main issues identified in 2012 leading the beam dumps were related to: Beam measurement quality; Front-end or software infrastructure problems; insufficient loop stability margin caused by the tighter loop constraints and requirements compared to 2010/11 operation. Many of these issues have been already addressed during 2012/13 and an important set of improvements are planned for LS1, notably the temperature controlled BPM racks and new Diode-Orbit acquisition system for the BPMs in the IR, ongoing improvements of the service infrastructure (FESA, CMW, TechNet, etc.), and upgrades of the OFC/OFSU infrastructure as outlined and confirmed by the FB review.

Better diagnostics, warning and status indication of the overall infrastructure, and better tracking and finer granularity of the given error assessment would be useful and help to track, understand and mitigate systematic, and rare but recurring errors of the feedbacks.

ACKNOWLEDGMENTS

The authors would like to thank the many colleagues at CERN and other labs who have been participating or contributed to the success of the LHC feedback systems: J. Andersson, M. Anderson, A. Boccardi, T. Bohl, A. Butterworth, E. Calvo Giraldo, P. Cameron (BNL), R. Denz, K. Fuchsberger, M. Gasior, S. Jackson, L. Jensen, R. Jones, M. Jonker, J.M. Jouanigot, K. Kasinski, Q. King, K. Kostro, M. Lamont, T. Lefevre, S. Page, L. Ponce, V. Ranjbar (FNAL), G. Sivatskiy, C.-Y. Tan (FNAL), E. Tedesco, J. Tückmantel, A. Verweij, J. Wenninger, W. Venturini, T. Wijnands, M. Zuin, and in particular our colleagues in the synchrotron-light-source community M. Böge, G. Decker, G. Rehm, and T. Schilcher.

REFERENCES

[1] T. Wijnands et al., "Requirements for Real Time Correction of Decay and Snapback in the LHC Superconducting Magnets", LHC-Project-Note-221 & Proceedings of EPCAC'00, 2000

[2] R. J. Steinhagen, 3rd Int. Workshop on Beam Orbit Stabilization ,IWBS'04, Grindelwald, Switzerland, 2004

[3] R. J. Steinhagen et al., "LHC Orbit Feedback Tests at the SPS", CERN-LHC-Project-Report-779, 2004

[4] Steinhagen et al., "LHC Orbit Stabilization Tests at the SPS", CERN-AB-2005-052, 2005

[5] R.J. Steinhagen, "LHC Beam Stability and Feedback Control - Orbit and Energy", CERN-AB-2007-049, 2007

[6] P. Cameron et al., "Advances Towards the Measurement and Control of LHC Tune and Chromaticity", DIPACâ05.

[7] R. Jones et al., "Towards a Robust Phase Locked Loop Tune Feedback System", DIPACâ05.

[8] Gasior et al., "An overview of the LHC Transverse Diagnostics Systems", LHC-PROJECT-Report-1166, 2008

[9] Steinhagen et al., "Advancements in the Base-Band-Tune and Chromaticity Instrumentation and Diagnostics Systems during LHC's First Year of Operation", CERN-BE-2011-016, 2011

[10] Steinhagen et al., "LHC Beam Stability and Performance of the Q/Q' Diagnostic Instrumentation", CERN-BE-2010-008, 2010

[11] Steinhagen et al., "Commissioning and Initial Performance of the LHC Beam-Based Feedback Systems", CERN-BE-2010-023, 2010

[12] R. J. Steinhagen, "Real-Time Beam Control at the LHC", CERN-BE-2011-020, PAC'11, NY, USA, 2011

[13] L. Ponce, "Tune and Orbit Feedback Performance: A User Perspective", LHC Beam Operation workshop – Evian 2011, Evian, France, 2011

[14] M. Gasior et al., "High Resolution Beam Orbit Measurement Electronics based on Compensated Diode Detectors", BIW'10, CERN-BE-2010-009, 2010

[15] M. Gasior et al., "A High-resolution Diode-Based Orbit Measurement System - Prototype Results from the LHC", DIPAC'11, CERN-BE-2011-024, 2011

[16] W. Höfle and D. Valuch, "Integration of ADT as alternate Q/Q' diagnostics source", private communication, 2012/2013

[17] J. Wenninger, "Feed-forward based on a multiple fill-to-fill averages of recurring FB trims", private communication, 2013

[18] M. Lamont et al., "Reviewers Recommendations", Architectural Review of the LHC Orbit & Tune Feedback Systems, `https://indico.cern.ch/conferenceDisplay.py?confId=241299`, CERN, 7th May, 2013