

Longitudinal emittance blowup in the large hadron collider

P. Baudrenghien¹, T. Mastoridis^{*}

BE-RF, CERN, Geneva, Switzerland

A B S T R A C T

The Large Hadron Collider (LHC) relies on Landau damping for longitudinal stability. To avoid decreasing the stability margin at high energy, the longitudinal emittance must be continuously increased during the acceleration ramp. Longitudinal blowup provides the required emittance growth. The method was implemented through the summer of 2010. Band-limited RF phase-noise is injected in the main accelerating cavities during the whole ramp of about 11 min. Synchrotron frequencies change along the energy ramp, but the digitally created noise tracks the frequency change. The position of the noise-band, relative to the nominal synchrotron frequency, and the bandwidth of the spectrum are set by pre-defined constants, making the diffusion stop at the edges of the demanded distribution. The noise amplitude is controlled by feedback using the measurement of the average bunch length. This algorithm reproducibly achieves the programmed bunch length of about 1.2 ns², at flat top with low bunch-to-bunch scatter and provides a stable beam for physics coast. The noise can be injected either in the beam phase loop or directly in the cavity voltage set point. These two different technical implementations are presented and their respective advantages analyzed. The performance of the algorithm and its further applications are also presented in this paper.

1. Introduction

The longitudinal emittance blowup is necessary to achieve the desired levels of Landau damping for longitudinal stability in the LHC. This fundamental motivation for the system development is presented in Section 2. A description of the algorithm, the choices of excitation noise spectrum, and the actual technical implementation are presented in Section 3. Section 4 presents the original blowup implementation through the LHC beam phase loop, its performance and limitations. Section 5 presents an alternative implementation through the LHC cavity controller, which allows for further applications of the longitudinal emittance blowup, especially selective excitation along the ring. Finally, Section 6 suggests possible future improvements for the system.

This paper follows the work previously presented by the authors in Ref. [1]. The LHC stability thresholds have been estimated in Ref. [2]. A theoretical treatment of the beam diffusion in the LHC has been presented in Ref. [3]. Using this treatment, a comparison of the estimated and measured longitudinal emittance growth rates was presented in Ref. [4].

2. Motivation for blowup

The first attempt to ramp single bunch, close to nominal intensity ($\approx 1.1 \times 10^{11}$ protons) took place on May 15th, 2010. At injection, the bunch was 1.2–1.3 ns long with about 0.4 eVs longitudinal emittance³ and this emittance was preserved during capture. Ramping was done with a constant 8 MV. Towards the end of the ramp, as the bunch length shrank down below 600 ps, a violent longitudinal instability developed as seen in Fig. 1, due to loss of Landau damping [2]. This behavior did not come as a surprise; it was consistent with LHC longitudinal stability studies done in 2000 [5]. At the time of the LHC design, the options of emittance blowup or an active longitudinal feedback system were considered to mitigate these instabilities [6]. The former solution was chosen for the LHC. The alternative option of distorting the longitudinal profile was not considered, but in this case periodic action on the bunch would be necessary during the long LHC coast, with implications for the physics program.

During acceleration, the threshold for loss of Landau damping scales as

$$\frac{\text{Im}(Z^{\text{thr}})}{n} < \frac{|\eta|E}{el_b\beta^2} \left(\frac{\Delta E}{E}\right)^2 \frac{\Delta\Omega_s}{\Omega_s} f_o \tau \quad (1)$$

¹ Tel.: +41 22 767 5581

^{*} Corresponding author. Tel.: +41 764 8790 68.

E-mail addresses: Philippe.Baudrenghien@cern.ch (P. Baudrenghien), tmastoridis@gmail.com (T. Mastoridis).

² Bunch length is defined as the 4σ value in this paper.

³ At CERN it is customary to quote the longitudinal emittance as $4\pi\sigma_z\sigma_E$. Note that, for a Gaussian distribution, 95% of the particles are within a $6\pi\sigma_z\sigma_E$ area. The $4\pi\sigma_z\sigma_E$ area contains 86.5% of the particles.




Fig. 1. First attempt to ramp nominal intensity single bunch. Bunch length during ramp. The longitudinal emittance is too low (< 0.4 eVs). The bunch becomes unstable. The bunch length measurement implies an oscillation with quadrupole components.

where $\text{Im}(Z^{thr})/n$ is the inductive impedance divided by the azimuthal bunch shape mode number n ($n=1$ dipole, $n=2$ quadrupole, etc.), η is the slip factor, E is the particle energy, e is the proton charge, I_b is the bunch current, $\beta = v/c$ with v the particle speed, Ω_s is the synchrotron frequency, f_o is the revolution frequency, τ is the bunch length, and $\Delta\Omega_s$ is the synchrotron tune spread [2]. The LHC values for these parameters are available in the Appendix.

Eq. (1) can be rearranged to

$$\frac{\text{Im}(Z^{thr})}{n} \propto \frac{\epsilon^{5/2}}{E^{5/4} V^{1/4}} \quad (2)$$

where ϵ is the longitudinal emittance and V is the total RF voltage. Since the LHC is always well above transition, η is approximately constant.

For a constant emittance the threshold quickly drops with energy, explaining the instability observed in the first ramp. The energy for the observed onset of instability is consistent with the 0.06Ω estimate for the inductive impedance divided by mode n for the LHC [2]. Since the bunches are stable at 450 GeV, it is sufficient to sustain a constant threshold to achieve stability throughout the LHC cycle, assuming that the longitudinal impedance is only marginally increased due to the collimator motion closer to the beam with energy increase. By inspection of Eq. (2), the stability margin is preserved if the emittance grows according to

$$\epsilon \propto E^{1/2} V^{1/10} \quad (3)$$

In the operational LHC blowup implementation, the bunch length τ is kept constant during the ramp. The emittance then grows as the bucket area (the bucket filling factor is constant)

$$\epsilon \propto E^{1/2} V^{1/2} \quad (4)$$

As the voltage increases during the ramp, the fixed bunch length blowup actually improves the stability margin during the acceleration.

The narrow-band impedance threshold was also studied in the RF design [5]. It is shown that, to avoid decreasing the threshold during the cycle, the emittance should be increased with energy at least as

$$\epsilon \propto \frac{E^{1/2}}{V^{1/6}} \quad (5)$$

Again the constant bunch length blowup results in a faster than strictly necessary emittance increase.

3. Longitudinal emittance blowup

The LHC blowup is inspired by the SPS system [7] but the LHC case is different: the much longer ramp makes the process smoother, there are short bunches in a single RF system with small synchrotron frequency spread, and there is almost no effect of bunch intensity (lower machine inductive impedance and much better compensation of the periodic beam loading). The beam is excited with RF phase noise acting via the fundamental RF system (400.8 MHz). The frequency of a single-particle synchrotron oscillation depends on the peak amplitude of its trajectory ϕ_{pk}

$$\Omega_s(\phi_{pk}) \approx \Omega_{s0} \left[1 - \left(\frac{\phi_{pk}}{4} \right)^2 \right] \quad (6)$$

with Ω_{s0} the synchrotron frequency of the zero-amplitude oscillation (Fig. 2).

This dependance can be used to selectively excite the particles in a chosen region centered around the core of the bunch. Assume, for example, that the phase noise spectrum extends from Ω_{s0} down to $0.85\Omega_{s0}$ (corresponding to an amplitude of phase oscillation equal to $\pi/2$ in Fig. 2). By exciting with a phase noise spectrum extending between these frequencies, the particles of the core of the bunch are driven in synchrotron resonance, but when the amplitude of their oscillation exceeds $\pi/2$, they would see no more coherent excitation. Diffusion should therefore stabilize around that point. The bunch length can be precisely controlled by fine adjustment of the lower frequency of the phase noise spectrum. For 1.2 ns target bunch length, the excitation is used in the band

$$0.85\Omega_{s0} \leq \Omega \leq 1.1\Omega_{s0} \quad (7)$$

The upper frequency exceeds Ω_{s0} to guarantee that the core is not missed (the filling factor is sufficiently low that higher modes are not excited). The beam intensity has a negligible impact on the incoherent synchrotron frequency shift in the LHC: the broadband inductive impedance ($\text{Im}(Z^{thr})/n \approx 0.06 \Omega$) reduces Ω_{s0} by only 1% at maximum bunch intensity, and the periodic beam loading is well below 0.5% in voltage [8]. A flat power spectral density is used. The excitation is applied during the acceleration ramp and the spectrum of the phase noise tracks the changing Ω_{s0} (Fig. 3). An algorithm has been developed for the generation of the phase noise samples with the required time-varying spectrum [9].

3.1. Feedback from measured length

When blowup was first tested in the LHC the bunch would indeed grow quickly till it reached the length corresponding to the lower synchrotron frequency in the excitation spectrum, but diffusion would not come to a complete stop then. The rate would just be reduced. An on-line measurement of the bunch length was




Fig. 2. Ω_s/Ω_{s0} as a function of the maximum phase deviation in radians. Stationary bucket.




Fig. 3. Synchrotron frequency in Hz, lower and upper frequencies for phase noise excitation during the ramp according to Eq. (7).

available from the LHC beam quality monitoring system [10] and could be used for feedback, to continuously adjust the amplitude of the noise, for a more precise control of the blowup. The algorithm updates the amplitude of the phase noise excitation x_n from measurement of the instantaneous 4- σ bunch length L_n (in radians, averaged over all bunches of one ring) and comparison with the target L_0

$$x_{n+1} = \alpha x_n + g(L_0 - L_n)$$

$$\text{if } x_{n+1} \leq 0 \quad \text{then } x_{n+1} \rightarrow 0$$

$$\text{if } x_{n+1} \geq 1 \quad \text{then } x_{n+1} \rightarrow 1 \quad (8)$$

Here n is the time index. One update is available every 5 s, limited by the rate of the beam quality monitor output (0.2 Hz). The variable x_n is a dimensionless factor (ranging from 0 to 1): the phase noise excitation signal is the product of x_n times a fixed level, corresponding to the maximum noise amplitude, presently set at 2° rms. For this reason, x_n is called the blowup gain. The algorithm is a simple low-pass filter, driven by the difference between measured length and target, with clamping. The excitation is switched off ($x=0$) if the length exceeds the target (bunch longer than desired) and it saturates at the maximum 2° rms ($x=1$). The parameter α defines the filtering of the beam quality monitor data, intended to reduce the measurement noise: $\alpha = 0.64$ is used, corresponding to a low-pass filter time constant of 11.2 s, i. e. an averaging over two beam quality monitor data points only. For good tracking α should be set for a reaction at least as fast as the observed bunch length transients. Its optimization has been empirical. For short bunches, lengthening caused by phase noise is proportional to the power spectral density $S_{\phi\phi}(f)$ sampled by the beam at the synchrotron frequency

$$\frac{dL^2}{dt} = \frac{8\Omega_{s0}^2}{\pi} S_{\phi\phi}\left(\frac{\Omega_{s0}}{2\pi}\right) \quad (9)$$

For a fixed noise level, the diffusion is fast at the beginning of the ramp (large synchrotron frequency) and tends to slow down with time as the synchrotron frequency decreases (Fig. 3). The parameter g is therefore a function scaled with Ω_{s0} and as such, it increases four-fold from the beginning to the end of the ramp to keep the effect on beam diffusion constant. For a good tracking (minimizing the deviations during the ramp) and a good precision (reproducibility of the end-ramp figure), it should be as large as possible. But its range is limited by stability considerations due to the low 0.2 Hz update rate.

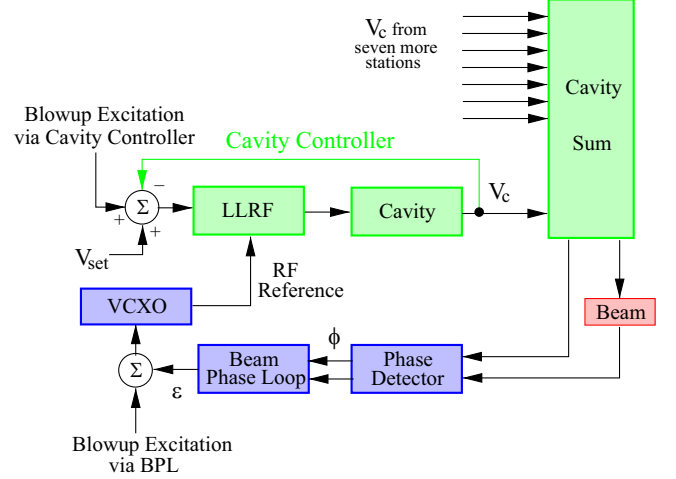


Fig. 4. Simplified block diagram of cavity controller and beam phase loop.

3.2. Technical implementation

The LHC RF includes many feedback control loops. Two essential loops for this analysis are the cavity controller loop (low level RF system) and the beam phase loop. Fig. 4 shows a simplified block diagram of the cavity controller and beam phase loop.

The cavity controller loop keeps the cavity voltage constant and reduces the RF station impedance experienced by the beam. It samples the RF field at 40.08 MHz and tries to control it for each bunch.

The beam phase loop input is the phase difference between the bunch and the total accelerating voltage as seen by the bunch. The beam phase loop averages this signal over all bunches of one ring and makes one correction per turn on the 400.8 MHz RF reference, to damp out barycentric longitudinal motion around the synchronous phase, motion driven by noise in the RF system or by other mechanisms.

The longitudinal blowup noise excitation can be added to the reference point of either one of these loops. As the beam phase loop updates at the revolution frequency f_{rev} , it applies the same noise on all bunches and can excite dipole mode zero only (excitation at $f_{RF} \pm f_s$). Acting through the cavity controller it is possible to apply different noise samples on successive bunches and excite on the revolution frequency sidebands (at $f_{RF} \pm n f_{rev} \pm f_s$).

The initial blowup implementation during 2011 employed the beam phase loop and is presented in Section 4. After extensive bunch length, beam and cavity field spectrum measurements (Section 4.2), moving the blowup to the cavity controller loop was investigated in the beginning of 2012 (Section 5).

4. Blowup through the beam phase loop

The phase noise is injected on the synchrotron side-bands of the RF frequency. If it were injected in the cavity drive directly, the beam phase loop, responsible for minimizing the noise in this very sensitive frequency band would cancel it [8]. The noise is therefore added as an offset in the beam phase loop. This results in the desired phase noise spectrum, between the beam phase (averaged over all bunches) and the cavity field but gives no direct control of the actual voltage, with some negative consequences described later in this paper.

Fig. 5 shows the performance of the blowup during a fill with 1380 nominal intensity bunches per ring. Displayed are the mean bunch length (averaged over the 1380 bunches of one ring) and

the instantaneous excitation level (blowup gain) during the 11 min long ramp (starting at minute 17, ending at 28 on the horizontal axis).

The target bunch length L_0 is set at 1.25 ns. Just before starting the ramp, the mean bunch length is 1.27 ns, for both beams. The adiabatic bunch shortening is clearly visible as soon as the ramp starts. The blowup feedback reacts and stabilizes the length after about 1 min. The following evolution is somewhat chaotic (notice the very fast jumps by more than 100 ps), but the algorithm correctly adapts the excitation level, reducing it when the bunch lengthens, and increasing it when it shrinks. Blow-up stops at the end of the ramp with, in this example, an achieved 1.18 ns in Beam 1 and 1.15 ns in Beam 2. The performance shown is typical: the fill to fill reproducibility is within ± 50 ps.

4.1. Bunch length equalization

A very good feature of the blowup is the reduction of the dispersion in bunch length: a ± 200 ps variation between the bunches is typical at the end of the injection plateau. Part of this spread is caused by the injector, the rest is due to intrabeam scattering, violent at injection energy, that blows-up the emittance of the bunches injected at the beginning of the filling sequence, which is never shorter than 15 min. After blowup in the LHC ramp the spread is reduced to ± 30 ps. Thanks to the band-limited phase noise spectrum, diffusion does indeed slow down at the desired amplitude. Figs. 5 and 6 correspond to the same fill. Fig. 6 shows the bunch length statistics, over the 1380 bunches of Beam 2, through the acceleration ramp: the overall ± 200 ps spread observed at the start of the ramp is reduced to ± 30 ps on flat top. The standard deviation is reduced from 60 ps to 15 ps.

4.2. Beam and cavity spectra with blowup through the beam phase loop

Heating of some machine components (beam screen, injection kickers, collimator jaws) was observed in 2011, with a strong correlation with bunch length [11,12]. This effect was unexpected based on the design simulations of these structures and their interaction with the beam current spectrum. The observed heating motivated detailed measurements of the average bunch profile and beam current spectrum.

There are two longitudinal wideband (3 GHz bandwidth) pick-ups per beam. They are used for the beam quality monitor, beam phase loop, and other purposes. One of them was used for the




Fig. 5. Bunch length (mean over 1380 bunches/beam) and excitation (blowup gain) during the ramp.




Fig. 6. Statistics on bunch length (mean, min, max and standard deviation errorbars) during the ramp. Beam 2, LHC Fill 2028.




Fig. 7. Single-sided beam spectrum. End of filling. Protons, 50 ns spacing.

beam spectra measurements. An effort has been made in the post-processing to correct for the dispersion due to ≈ 380 m cable to the surface, which dominates the signal distortion on this chain. The 50 ns bunch spacing leads to 20 MHz sampling lines. The non-uniform filling pattern causes the line-to-line amplitude variation. The average bunch spectrum can be deduced by the envelope of this measurement. Finally, the IF bandwidth for this measurement is 20 kHz, which results in averaging over a fraction of a turn ($f_{rev} \approx 11$ kHz).

Fig. 7 shows the beam spectrum at the end of the filling, for a proton beam. There is significant power around 1.6 GHz, about 35 dB below the 400 MHz component (factor of 3000 in power). This measurement implies that the longitudinal distribution of the LHC bunches is non-Gaussian, already at the injection from the SPS (a Gaussian profile transforms into a smooth parabolic spectrum in logarithmic scale). This shape is produced by the longitudinal blowup during the acceleration ramp in the SPS [13].

Fig. 8 shows the beam spectrum at the end of the ramp, for the same LHC fill. The high frequency side lobe gets amplified during the ramp to ≈ 30 dB below the 400 MHz component, so the longitudinal blowup seems to further distort the longitudinal distribution. Studies are in progress from the LHC impedance team to determine possible overlap of this side lobe with resonances of the structures that exhibit sensitivity to heating. It should be noted that this side lobe is present during the physics coast with a very slowly reducing amplitude. Even after 10 h it is only reduced by 5–10 dB.




Fig. 8. Single-sided beam spectrum. End of ramp. Protons, 50 ns spacing.




Fig. 9. Average longitudinal bunch distribution, reconstructed from beam spectrum. End of filling. Protons, 50 ns spacing.

The time domain signal can be reconstructed using an inverse Fourier transform to extract the average longitudinal distribution. The beam spectra measurements are done with a spectrum analyzer and do not provide phase information. At injection and in collision though, the stable phase is 180° and the bucket is symmetric, therefore resulting in a symmetric bunch profile. As a result, the phase of its Fourier transform is 0 or 180° . To convert the beam spectra to the time domain, a 180° phase shift has been assumed for the first side lobe (above 1.3 GHz). It should be mentioned that the phase dispersion due to the cable corrugation becomes more significant in the time domain case. Again, a correction has been applied in post-processing.

Fig. 9 shows the reconstructed longitudinal bunch distribution at the end of filling (450 GeV) compared to a Gaussian profile of same peak and area. As expected, the bunch spectrum is not Gaussian, but it is closer to a Gaussian with finite tails distribution [14] (denser core, lower/finite tails). Compared to a Gaussian shape, this distribution reduces capture losses in the LHC. At the end of the ramp, the distribution has distorted further (Fig. 10). Such a distribution is beneficial for longitudinal stability considerations [15] and longitudinal lifetime (depleted tails), but it might contribute to increased machine heating. It is especially interesting to see the area around 600 ps. The edge of the blowup excitation ($0.85\Omega_{s0}$) is exactly at 600 ps. The longitudinal distribution implies that the excitation is stronger at the edge of the blowup range rather than at the core of the beam.




Fig. 10. Average longitudinal bunch distribution, reconstructed from beam spectrum. End of ramp. Protons, 50 ns spacing.




Fig. 11. Single-sided beam spectrum. End of filling. Ions, 200 ns spacing.

Observations of ion beam spectra provided similar results. The low intensity ion bunch is stable without longitudinal blowup. However, the increased longitudinal emittance leads to smaller transverse emittance growth driven by intrabeam scattering and therefore higher luminosity. Ion bunches are spaced by 200 ns, therefore the frequency spectrum consists of lines 5 MHz apart (the proton beam spectra presented in this work consist of 20 MHz lines).

Due to the diffusion from the higher intrabeam scattering rates for ion bunches, the bunch distribution returns to a Gaussian relatively quickly after a disturbance. As such, it is easier to distinguish the effects of the blowup. For ions, the spectrum at the start of the ramp is Gaussian, with no visible side lobe, after the diffusion taking place in both SPS and LHC flat bottom (Figs. 11 and 12). When longitudinal blow-up starts in the LHC ions ramp to maintain the desired 1.2 ns bunch length, the spectrum shows higher frequency components and the profile differs much from Gaussian (Figs. 13 and 14). The spectrum shows a side lobe extending from 1.3 GHz to 2.6 GHz. After a few hours into physics (2 to 3), the ion beam spectrum returns to a Gaussian shape, as consequence of intrabeam scattering⁴ (Figs. 15 and 16). This clearly

⁴ The longitudinal emittance growth rate due to intrabeam scattering is about 6 h for the beam parameters of the 2011 LHC ion run.




Fig. 12. Average longitudinal bunch distribution, reconstructed from beam spectrum. End of filling. Ions, 200 ns spacing.




Fig. 14. Average longitudinal bunch distribution, reconstructed from beam spectrum. End of ramp. Ions, 200 ns spacing.




Fig. 13. Single-sided beam spectrum. End of ramp. Ions, 200 ns spacing.




Fig. 15. Single-sided beam spectrum. 3.3 h into physics. Ions, 200 ns spacing.

identified the link between bunch profile and blowup scheme, and motivated the developments of alternative methods in 2012.

4.3. Cavity voltage spectra with blowup through the beam phase loop

These measurements prompted a closer inspection of the cavity phase noise spectrum during the ramp. Fig. 17 shows the phase noise of the cavity sum for Beam 1 (eight cavities per beam), right before the end of the ramp with the blowup on. It is obvious from this image that the achieved noise spectrum is not flat as desired, but has a notch close to the synchrotron frequency. This results in the non-Gaussian longitudinal distribution after the ramp, as confirmed by a diffusion model to be presented in an upcoming publication. The action of the beam phase loop amplified by the high gain of the beam response at the synchrotron resonance introduces the notch in the synchrotron excitation, as explained below.

4.4. Simple beam phase loop model

Following the simplified block diagram in Fig. 18 and ignoring the cavity response, it can be shown that the cavity phase noise




Fig. 16. Average longitudinal bunch distribution, reconstructed from beam spectrum. 3.3 h into physics. Ions, 200 ns spacing.




Fig. 17. Cavity sum phase noise. Beam 1.




Fig. 18. Simplified block diagram of beam phase loop.




Fig. 19. Estimated cavity phase noise power spectral density.

power spectral density $S_{cav}(f)$ is

$$S_{cav}(f) = \left| \frac{G}{s + G(1 - H_{beam}(s))} \right|^2 N(f) \quad (10)$$

where $N(f)$ is the power spectral density of the injected noise and $H_{beam}(s)$ is the beam response. For this simple model the beam response has been approximated by a damped oscillator at the synchrotron frequency f_s with very high damping time.

The resulting cavity phase noise spectral density can then be estimated as a function of the intended excitation, as shown in Fig. 19. This spectrum shows a clear correlation with that in Fig. 17, implying that the distortion in the cavity noise spectrum and longitudinal distribution are a result of the action of the beam phase loop on the excitation through the beam. This is a limitation of the technical implementation through the beam phase loop, which reduces the efficiency of the longitudinal emittance blowup.

5. Blowup through the cavity controller

The above discovery motivated the investigation of an alternative technical implementation, effectively injecting the blowup excitation through the cavity voltage set point in the cavity controller. If this excitation is on the synchrotron frequency sideband, the beam phase loop would register an ensemble motion of all bunches and would damp it to zero. However, if the excitation is introduced on the synchrotron sideband of a revolution harmonic and with a symmetric filling pattern, the average phase deviation over a turn – and consequently the beam phase loop error signal – is zero. Fig. 20 shows the injected noise power spectral density with the new implementation. An early test of the new algorithm at 3.5 TeV showed promising results in achieving a much smoother bunch lengthening (Fig. 21), significantly lower high frequency components and a longitudinal distribution approaching a Gaussian (Figs. 22 and 23).

Such an implementation has the additional advantage of allowing for selective excitation during a turn. This technique has already been used for various machine development studies in the LHC in 2011. For example, it was used to achieve a wide spread of bunch lengths (1.25–1.7 ns) for a fill with eight equidistant bunches of common initial bunch lengths of 1 ns, as shown in Fig. 24. The new longitudinal emittance blowup was tested during




Fig. 20. Injected noise power spectral density around $f_{rev} \pm f_s$. Horizontal axis shifted by $f_{RF} + f_{rev}$. ($f_{rev} = 11.245$ kHz).




Fig. 21. Bunch lengthening with blowup through the cavity controller at 3.5 TeV. Ion Beam.




Fig. 22. Single-sided beam spectrum. Longitudinal blowup through the cavity controller. 3.5 TeV, ion beam.




Fig. 23. Average longitudinal bunch distribution, reconstructed from beam spectrum. Longitudinal blowup through the cavity controller. 3.5 TeV, ion beam.

the 2012 LHC commissioning, but it has not been used during normal operation yet.

5.1. Batch-by-batch blowup

This new blowup implementation could be used for batch-by-batch blowup at LHC injection. The transverse emittance in the LHC increases by 20% from 450 GeV to 4 TeV, mostly due to intrabeam scattering. By rapidly increasing the longitudinal emittance of each newly injected batch, the intrabeam scattering growth rates for the transverse emittance are reduced, thus achieving lower transverse emittance in physics and higher luminosity.

A machine development study was conducted on April 22nd 2012 [16] to test the feasibility and performance of the batch-by-batch longitudinal emittance blowup. The selective blowup of longitudinal emittance worked very well. It was possible to excite individual batches spaced by 925 μ s (nominal batch spacing in the LHC) without affecting the leading and trailing (if any) batches, in a very short time period. Even though the cavity bandwidth is only about ± 10 kHz ($Q_L=20\,000$), with the strong RF feedback the field controls bandwidth is ± 350 kHz resulting in a 10–90% rise time of ≈ 1 μ s.




Fig. 24. Phase noise excitation, applied at 3.5 TeV, with a spectrum centered on the first revolution frequency sideband, with a sine amplitude along the ring, resulting in different lengths for the eight circulating bunches. Machine development session on August 24th 2011.




Fig. 25. Average bunch length per batch. Fill 2556, Beam 2. Eleven batches of 144 bunches each were injected, separated by 925 ns. Batch-by-batch blowup active for batches 3–11, with varied target lengths and excitation amplitudes.

Fig. 25 shows the average bunch length per batch for fill 2556, Beam 2. Eleven batches of 144 bunches each were injected. Initially, the bunches come in with a bunch length of about 1.5 ns, as measured at SPS extraction. In a few milliseconds though, the bunch length is reduced to about 1.1 ns due to the mismatched capture voltage. After that, the bunch length increases with about 400 ps/h due to intrabeam scattering. During this injection, the batch-by-batch blowup was off for the first two batches and then turned on for the rest. Within a couple of minutes the excitation has increased the bunch length to the target value (varied among batches for machine development purposes – the feedback on the amplitude of the phase noise excitation can be set independently for each batch). The maximum excitation amplitude was reduced by a factor of 2 for the last batch leading to smoother – but slower – approach to the target value. Notice that there is no effect in the circulating batches when blowup is applied on the newly injected ones (still spaced by 925 ns).

From the transverse emittance data analysis, there is some early indication that growth rates are indeed lower for blown up batches. Data from physics fills once the system is operational will provide more accurate results, since there will be more bunches

per batch and bunch-by-bunch specific luminosity data will be available too. The system was commissioned in late 2012. Conclusions on any performance improvements achieved are still pending.

6. Possible improvements

It should be easy to improve the precision of the blowup by increasing the gain g of the feedback algorithm (Eq. (8)). Unfortunately, loop stability sets the limit: when the reaction time gets close to the latency between measurements, a sampled feedback system will oscillate. An upgrade of the beam quality monitor is therefore underway, to increase the data rate.

A more fundamental limitation comes from the definition of bunch length in the presence of non-adiabatic change of bunch shape. In Fig. 5, at time 22 min, the Beam 2 mean bunch length jumps by 150 ps in only 30 s. Such a fast reaction is not physically possible without a change of bunch profile. The beam quality monitor extracts the full width at half maximum for each bunch, and estimates the 4σ equivalent length assuming a Gaussian profile. Rapid changes of beam profile have been observed during the ramp, which have a big impact on the full width at half maximum measurement and result in the observed transients. Fig. 26 shows a comparison of the full width at half maximum estimate and a Gaussian fit during the ramp. It is clear from this image that the distribution (and its relationship to a Gaussian) changes significantly during the ramp. During rapid changes of profile, it is not clear how any measurement can precisely drive the amplitude of the blowup noise. A study of possible correlation of these fast transients with the distribution of bunch lengths along the ring, the mean bunch length at the beginning of the ramp, or the bunch intensities, was unsuccessful so far. Another tentative explanation for these jumps is the crossing of the 50 Hz synchrotron frequency line during the ramp but this happens less than 2 min into the ramp, much earlier than the observed jumps in the bunch length.

The longitudinal blowup could be used in the future to distinguish the heating effects from each beam separately, by varying the blowup settings for each beam independently.

Additionally, it should be possible to study the RF heating as a function of the bunch profile, to identify frequencies of interest in the offending impedances. Some early measurements were conducted in 2012 [17].




Fig. 26. Comparison of the full width at half maximum estimate and Gaussian fit during the ramp. Fill 1865.

7. Conclusions

The longitudinal emittance blowup is essential for the stability of the LHC proton beam. The bunch length is kept at a set-value during the ramp, thereby providing sufficient longitudinal emittance increase to preserve Landau damping. Stabilization of the bunch length is also essential to limit the beam induced heating of some machine elements (beam screen and kickers). The blowup also reduces the spread in bunch length during physics, improving the beam quality and its overall luminosity. In addition, it leads to a reduction of the transverse emittance growth due to intrabeam scattering, with further luminosity benefits.

Two technical implementations of the longitudinal blowup are discussed. The one through the beam phase loop has been in operation since 2010 and can be used for any filling pattern during the ramp, but it slightly distorts the longitudinal distribution. The implementation through the cavity controller allows for selective excitation, which is very useful for batch-by-batch blowup at injection, as well as for machine development purposes. It requires a symmetric filling pattern during the ramp, so it has not been used in normal operation during the ramp yet. It has been successfully tested for batch-by-batch blowup at injection [16].

Acknowledgements

T. Bohl, T. Linnecar, E. Shaposhnikova, and J. Tuckmantel developed and commissioned the SPS longitudinal emittance blowup, which provided useful background for the LHC implementation. The authors would also like to acknowledge J. Tuckmantel's significant contributions on the development of the noise generation algorithms, the simulations of the algorithm behavior in the LHC, and his advice in all stages of this work. E. Shaposhnikova designed the optimal shaping of the excitation phase noise spectrum used in the LHC and contributed in determining the limitations of the implementation through the beam phase loop. The authors would like to thank M. Jaussi and J. Molendijk for their many contributions on the software and firmware development necessary for the blowup algorithm. We would also like to thank G. Papotti for her contributions on the beam quality monitor measurements necessary for the bunch length feedback. Discussions with A. Butterworth contributed to the development of the technical implementation.

Appendix

Some of the design LHC parameters [18], as well as the 2011, 2012 operational parameters are shown here for reference (Table 1).

Table 1

LHC parameters at 450 GeV and 7 TeV (design report) and from the 2011 and 2012 LHC runs. τ is the 4-sigma bunch length.

E (GeV)	Design report		2011		2012	
	450	7000	450	3500	450	4000
η (10^{-4})	3.182	3.225	3.173	3.216	3.173	3.216
I_b DC (A)	0.582		0.38		0.41	
f_0 (kHz)			11.245			
f_{RF} (MHz)			400.8			
f_s (Hz)	61.8	21.4	55.1	28.1	55.1	24
V_{RF} (MV)	8	16	6	12	6	12
τ (ns)	1.5	1	1.5	1.2	1.5	1.3

References

- [1] P. Baudrenghien, et al., Longitudinal emittance blowup in the LHC, in: Proceedings of the IPAC 2011, 4–9 September 2011, San Sebastian, Spain.
- [2] E. Shaposhnikova, et al., Loss of Landau damping in the LHC, in: Proceedings of the IPAC 2011, 4–9 September 2011, San Sebastian, Spain.
- [3] T. Mastorides, et al., RF system models for the LHC with application to longitudinal dynamics, *Physical Review Special Topics: Accelerators and Beams* 13 (2010) 102801.
- [4] T. Mastoridis, et al., Radio frequency noise effects on the CERN large Hadron collider beam diffusion, *Physical Review Special Topics: Accelerators and Beams* 14 (2011) 092802.
- [5] E. Shaposhnikova, et al., Longitudinal beam parameters during acceleration in the LHC, LHC Project Note 242, December 2000.
- [6] D. Boussard, et al., Is a longitudinal feedback system required for LHC? LHC Project Note-205, November 1999.
- [7] J. Tuckmantel, et al., Study of controlled longitudinal emittance blow-up for high intensity LHC beams in the CERN SPS, in: Proceedings of the EPAC 2008, 23–27 June 2008, Genoa, Italy.
- [8] P. Baudrenghien, et al., The LHC RF system, experience with beam operation", in: Proceedings of the IPAC 2011, 4–9 September 2011, San Sebastian, Spain.
- [9] J. Tuckmantel, Digital generation of noise-signals with arbitrary constant or time-varying spectra, in: Proceedings of the EPAC 2008, 23–27 June 2008, Genoa, Italy.
- [10] G. Papotti, et al., Longitudinal beam measurements at the LHC: the LHC beam quality monitor, in: Proceedings of the IPAC 2011, 4–9 September 2011, San Sebastian, Spain.
- [11] B. Salvant, et al., Beam induced heating, in: Proceedings of the LHC Beam Operation Workshop 2011, 12–14 December 2011, Evian, France.
- [12] E. Metral, Beam-Induced Heating/Bunch Length/RF and Lessons from 2012, LHC Performance Workshop, 6–10 February 2012, Chamonix, France.
- [13] T. Bohl, et al., Controlled Longitudinal Emittance Blow-up in the SPS as injector and LHC Test-Bed, CERN-AB-Note-2003-084-MD, December 2003.
- [14] J.L. Laclare, Bunched Beam Coherent Instabilities, CERN 87-03, CAS on General Acc. Phys., Oxford, UK, 1985, pp. 264–326.
- [15] C.Y. Tan, A. Burov, Phase modulation of the bucket stops bunch oscillations at the fermilab tevatron, *Physical Review Special Topics: Accelerators and Beams* 15 (2012) 044401.
- [16] T. Mastoridis, et al., Batch by batch longitudinal emittance blowup MD, CERN-ATS-Note-2012-050 MD, June 2012.
- [17] B. Salvant, et al., Beam induced RF heating, in: Proceedings of the LHC Beam Operation Workshop 2012, 17–20 December 2012, Evian, France.
- [18] LHC Design Report, vol. 1, Chapter 2, CERN-2004-003, 4 June 2004.