

LONGITUDINAL EMITTANCE BLOW-UP IN THE LHC

P. Baudrenghien[#], A. Butterworth, M. Jaussi, T. Mastoridis, G. Papotti, E. Shaposhnikova, J. Tuckmantel, CERN, Geneva, Switzerland

Abstract

The 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 blow-up provides the required emittance growth. The method was implemented through the summer of 2010. We inject band-limited RF phase-noise in the main accelerating cavities during the whole ramp of about 11 minutes. 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 ns (4 σ) at flat top with low bunch-to-bunch scatter and provides a stable beam for physics coast.

MOTIVATION FOR BLOW-UP

The first attempt to ramp single bunch, nominal intensity ($1.1 \cdot 10^{11}$ p) took place on May 15th, 2010. At injection, the bunch was 1.2-1.3 ns long (4 σ), with 0.3-0.4 eVs longitudinal emittance 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, due to loss of Landau damping [1]. This behaviour did not come as a surprise; it was consistent with LHC longitudinal stability studies done in 2000 [2]. During acceleration, the threshold for loss of Landau damping scales as [1]

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

where ε is the longitudinal emittance. 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 estimated 0.06 Ω inductive impedance of the LHC [1]. The LHC RF design specified longitudinal blow-up during the ramp to keep the threshold constant [2]. By inspection of equation (1), the stability margin is preserved if the emittance grows according to

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

In the operational LHC blow-up implementation, we keep the bunch length L constant during the ramp. The

emittance then grows as the bucket area (the bucket filling factor is constant). We have

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

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

The narrow-band impedance threshold was also studied in the RF design [2]. It is shown that, to avoid decreasing the threshold during the cycle, the emittance should be increased with energy at least as

$$\varepsilon \propto E^{1/2} / V^{1/6} \quad (4)$$

Again the constant bunch length blow-up results in a faster than strictly necessary emittance increase.

LONGITUDINAL EMITTANCE BLOW-UP

The LHC blow-up is inspired by the SPS system [3] but the LHC case is different: much longer ramp making the process smoother, short bunches in a single RF system with small synchrotron frequency spread, and almost no effect of bunch intensity (lower machine inductive impedance and much better compensation of the periodic beam loading). We excite the beam with RF phase noise acting on 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}

$$\Omega_s(\phi_{pk}) \approx \Omega_{s0} \left[1 - \left(\frac{\phi_{pk}}{4} \right)^2 \right] \quad (5)$$

with Ω_{s0} the synchrotron frequency of the zero-amplitude oscillation (figure 1).




Figure 1: Ω_s/Ω_{s0} as a function of the maximum phase deviation in radian. Stationary bucket.

This dependance can be used to selectively excite the particles in a chosen region centred 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

[#]Philippe.Baudrenghien@cern.ch

to $\pi/2$ in Figure 1). By exciting with a phase noise spectrum extending between these frequencies, we would drive the particles of the core of the bunch 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 4 σ target bunch length, we use excitation in the band

$$\frac{6}{7}\Omega_{s0} \leq \Omega \leq 1.1\Omega_{s0} \quad (6)$$

The upper frequency exceeds Ω_{s0} to guarantee that we do not miss the core. The beam intensity has a negligible impact on the incoherent synchrotron frequency shift in the LHC: the broadband inductive impedance (0.06 Ω) reduces Ω_{s0} by 1% only at maximum bunch intensity, and the periodic beam loading is well below 0.5 % in voltage [4]. We use a flat Power Spectral Density (PSD). The excitation is applied during the acceleration ramp and the spectrum of the phase noise must track the changing Ω_{s0} (Figure 2).




Figure 2: Synchrotron frequency in Hz (blue), lower and upper frequencies for phase noise excitation during the ramp according to equation (6).

An algorithm has been developed for the generation of the phase noise samples with the required time-varying spectrum [5].

FEEDBACK FROM MEASURED LENGTH

When blow-up 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 available from the LHC Beam Quality Monitoring system (BQM) [6] and could be used for feedback, to continuously adjust the amplitude of the noise, for a more precise control of the blow-up. The algorithm updates the amplitude of the phase noise excitation x_n from measurement of the instantaneous bunch length L_n (averaged over all bunches of one ring) and comparison with the target L_0

$$\begin{aligned} x_{n+1} &= a \cdot x_n + g \cdot (L_0 - L_n) \\ \text{if } x_{n+1} \leq 0 \text{ then } x_{n+1} &\rightarrow 0 \\ \text{if } x_{n+1} \geq 1 \text{ then } x_{n+1} &\rightarrow 1 \end{aligned} \quad (7)$$

Here n is the time index. We have one update every 5 seconds, limited by the rate of the BQM 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 degrees rms. For this reason, x_n is called the *Blowup Gain*. The algorithm is a simple low-pass filter (LPF), 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 degrees rms ($x=1$). The parameter a defines the filtering of the BQM data, intended to reduce the measurement noise: we use $a=0.64$, corresponding to a LPF time constant of 11.2 seconds, i.e. an averaging over two BQM data points only. For good tracking a 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 PSD $S_{\phi\phi}(f)$ sampled by the beam at the synchrotron frequency

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

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 (Figure 2). The parameter g is therefore a function, increased 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 (see below). Figure 3 shows the performance of the blow-up 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 minutes long ramp (starting at minute 17, ending at 28 on the horizontal axis).




Figure 3: Bunch length (mean over 1380 bunches/beam) and excitation (Blowup Gain) during the ramp.

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 blow-up feedback reacts and stabilizes the length after about one minute. 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.

BUNCH LENGTH EQUALIZATION

A very good feature of the blow-up is the reduction of the dispersion in bunch length: at the end of the injection plateau we would typically have ± 200 ps variation between the bunches. Part of this spread is caused by the injector, the rest is due to the Intra Beam 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 minutes. After blow-up 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. Figures 3 and 4 correspond to the same fill. Figure 4 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.




Figure 4: Statistics on bunch length (mean, min, max and standard deviation errorbars) during the ramp.

IMPROVEMENTS

It should be easy to improve the precision of the blow-up by increasing the gain g of the feedback algorithm (7). Unfortunately we are limited by the loop stability: when the reaction time gets close to the latency between measurements, a sampled feedback system will oscillate. An upgrade of the BQM is therefore underway, to increase the data rate.

A more fundamental limitation comes from the definition of bunch length in presence of non-adiabatic change of bunch shape. In figure 3, at time 22 minutes, the beam 2

mean bunch length jumps by 150 ps in only 30 seconds. Such a fast reaction is not physically possible without a change of bunch profile. The BQM extracts the Full Width at Half Maximum (FWHM) for each bunch, and estimates the 4σ equivalent length assuming a Gaussian profile [6]. Rapid changes of beam profile have been observed during the ramp, which have a big impact on the FWHM measurement and result in the observed transients. During rapid changes of profile, it is not clear how any measurement can precisely drive the amplitude of the blow-up 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 was not confirmed by the observations. 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 (BPL), responsible for minimizing the noise in this very sensitive frequency band would cancel it [4]. The noise is therefore added as an offset in the BPL. 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. An alternative is to inject the noise on a revolution frequency harmonic (at $\omega_{rf} \pm n\omega_{rev} \pm \Omega_s$) that is invisible to the BPL with a symmetric machine filling as it averages over one turn.

CONCLUSIONS

Longitudinal blow-up is essential for the stability of the LHC beam. We keep the bunch length 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). In addition the blow-up reduces the spread in bunch length during physics, improving the beam quality and its overall luminosity.

REFERENCES

- [1] E. Shaposhnikova et al., Loss of Landau damping in the LHC, these proceedings
- [2] E. Shaposhnikova, Longitudinal beam parameters during acceleration in the LHC, LHC Project Note 242, Dec 2000
- [3] J. Tuckmantel et al., Study of Controlled Longitudinal Emittance Blow-up for High Intensity LHC beams in the CERN SPS, EPAC08
- [4] P. Baudrenghien et al., The LHC RF System, Experience with Beam Operation, these proceedings
- [5] J. Tuckmantel, Digital Generation of Noise-Signals with Arbitrary Constant or Time-Varying Spectra, LHC Project Report 1055, Feb. 2008
- [6] G. Papotti et al., Longitudinal Beam Measurements at the LHC: the LHC Beam Quality Monitor, these proceedings