

RUNNING THE RF AT HIGHER ENERGY AND INTENSITY

P. Baudrenghien, T. Mastoridis, CERN, Geneva, Switzerland

Abstract

The improvements done to the RF parameters and hardware in 2011 are reviewed. Then the upgrades planned for 2012 are presented: further reduction of capture losses with the longitudinal damper, batch by batch blow-up at injection and modification of the controlled blow-up to preserve bunch profile. Operation at higher energy is readily possible with the present RF power, and does not degrade longitudinal stability thanks to the controlled longitudinal emittance growth during the ramp. For operation with higher beam current, the observations in 2011 indicate that there is no single bunch instability issue with up to $3 \cdot 10^{11}$ p per bunch. With the large gain of the RF feedback and One-Turn feedback, the cavity impedance at the fundamental will not be a limitation for ultimate intensity ($1.7 \cdot 10^{11}$ p per bunch) with 25 ns spacing. The klystron power (300 kW RF at saturation) is sufficient for 25 ns operation with nominal intensity (2808 bunches per beam, $1.1 \cdot 10^{11}$ p per bunch). An RF roadmap for going beyond will be outlined: it calls for an upgrade of the LLRF only and should allow for operation with ultimate beam intensity (25 ns spacing, 2808 bunches, $1.7 \cdot 10^{11}$ p per bunch) after Long Shutdown one.

NEW FEATURES 2011 VERSUS 2010

The statistics on RF faults has been presented at the Evian workshop [1] and will not be repeated here.

Increased capture voltage

At extraction, the SPS RF is 7 MV (200 MHz). The bunch has 1.5 ns length* (4σ), and $4.5 \cdot 10^{-4}$ energy spread $\Delta E/E$ ($2\sigma_E$), resulting in a $4\pi\sigma_r\sigma_E$ emittance of 0.5 eVs^\dagger . In 2011 the LHC capture voltage was increased from 3.5 MV (2010) to 6 MV. The bucket area was increased from 0.9 eVs in 2010, to 1.2 eVs. The bucket half height is now $9.6 \cdot 10^{-4} \Delta E/E$. With 7 MV at 200 MHz, the SPS bucket area at extraction is 3 eVs but the longitudinal distribution is limited to a much smaller region: controlled longitudinal blow-up is applied during the SPS ramp to keep the beam stable [2]. The blow-up is turned off near the moment when the voltage program corresponds to a bucket area of 1.05 eVs only. The voltage is then raised adiabatically to 7 MV for bunch shortening before transfer to the LHC. We can therefore

* We quote the 4σ length of a Gaussian bunch having the same Full Width at Half Maximum as the measured bunch.

[†] At CERN it is customary to quote the longitudinal emittance as $4\pi\sigma_r\sigma_E$. Note that, for a Gaussian distribution, and small filling factor, 95% of the particles are within a $6\pi\sigma_r\sigma_E$ area. The $4\pi\sigma_r\sigma_E$ area contains 86.5% of the particles.

limit the SPS bunch to a 1.05 eVs contour in a stationary 7 MV bucket (at 200 MHz). Figure 1 shows the situation: the 1.05 eVs contour falls almost entirely within the LHC bucket. Assuming a Gaussian distribution for the SPS bunch, truncated at the 1.05 eVs contour, the calculated loss is 0.02%.

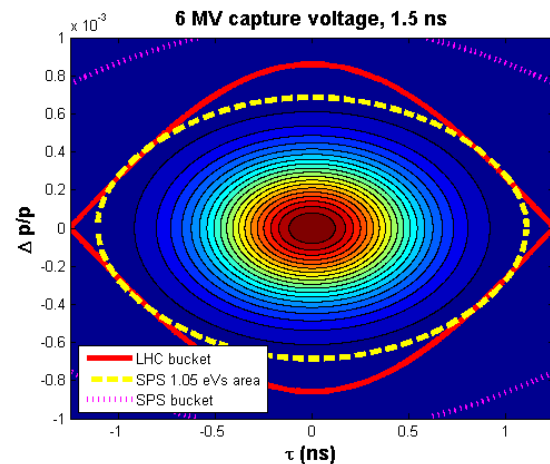


Figure 1: Longitudinal phase space at injection: We assume a Gaussian distribution for the SPS bunch and display contours corresponding to steps of 5% in integrated intensity. The Gaussian is truncated at the 1.05 eVs contour (yellow).

In figure 2 we introduce a small injection error (100 ps and $10^{-4} \Delta p/p$). This results in a small portion of the bunch falling outside the LHC bucket (calculated 0.4 % loss with the truncated Gaussian model). In 2011 we have observed 0.5% loss from injection to start ramp.

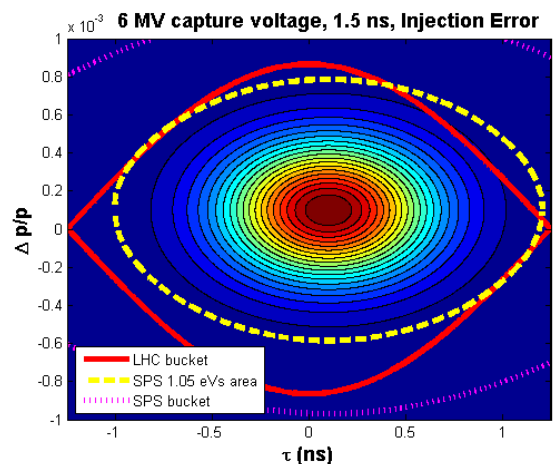


Figure 2: Longitudinal phase space at injection with a small error: 100 ps and $10^{-4} \Delta p/p$.

A consequence of the voltage mismatch (matched voltage is around 2.5 MV) is the bunch length reduction

after capture (from 1.5 ns to 1.1 ns). We could take advantage of the large available bucket to blow up the longitudinal emittance after each injection and restore the 1.5 ns length (batch-by-batch blowup). With 1.5 ns and 6 MV, we get 0.83 eVs ($4\pi\sigma_E\sigma_t$) emittance. We could increase it further by capturing with 8 MV as for the Lead ions, leading to 0.97 eVs. This is planned for 2012.

Larger voltage in physics

In 2011, controlled longitudinal emittance blow-up was applied in the eleven minutes long ramp, keeping the bunch length around 1.2 ns, while the RF voltage was increased linearly from 6 MV to 12 MV. (In 2010 we used 8 MV only in physics). The 12 MV provide a larger longitudinal emittance, thereby reducing the transverse emittance growth due to Intra Beam Scattering. At the beginning of the 3.5 TeV flat top we now have 2 eVs longitudinal emittance in a 4.7 eVs bucket (1.5 eVs in a 3.8 eVs bucket in 2010).

One Turn Feedback

The One Turn Feedback (OTFB) was commissioned on all cavities in October 2011. It produces gain around the revolution frequency sidebands only, thereby compensating for the transient beam loading and reducing the effective cavity impedance at the fundamental RF [3]. Figure 3 shows its performance with 0.38 A DC beam current, the highest achieved in 2011 (2100 bunches, 10^{11} p/bunch, 25 ns spacing). The transient beam loading caused by the gaps between the batches and by the 10 μ s long abort gap are clearly visible in the cavity field amplitude and phase. Shown are the voltage amplitude and phase in a cavity without OTFB (C7B1) and with OTFB (C1B1). With the One Turn Feedback, the voltage modulation is barely visible (0.3%) and the phase modulation is 0.5 degree pk-pk. Not shown is the demanded klystron power. The transients are actually reduced with the One Turn Feedback. Figure 4 shows the phase noise Power Spectral Density (PSD) of a cavity in physics conditions, but without beam. It compares the spectrum without and with OTFB. On the revolution frequency harmonics the feedback reduces the phase noise power. Although the PSD is increased halfway between revolution lines, the beam does not respond at these frequencies.

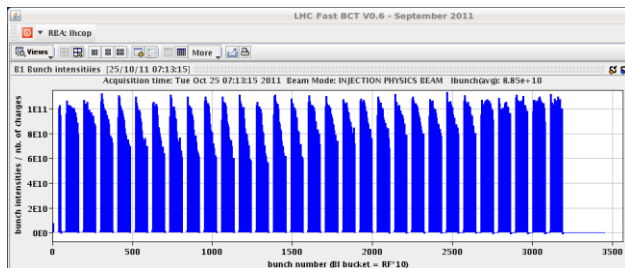


Figure 3a: Beam current. The very uneven bunch intensity is caused by the poor transverse lifetime with 25 ns spacing.

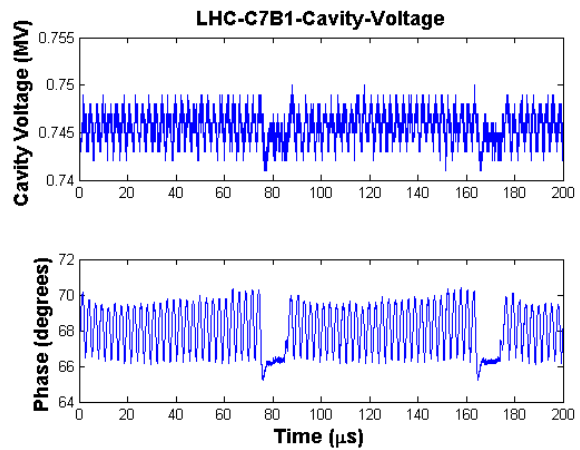


Figure 3b: Cavity 7 Beam 1 field voltage and phase. The OTFB is OFF.




Figure 3c: Cavity 1 Beam 1 field voltage and phase. The OTFB is ON.




Figure 4: Phase noise PSD of Cavity 1 Beam 1 in physics conditions (1.5 MV, $Q_L=60k$), from DC to 500 kHz, no beam. OTFB on (blue trace) and OTFB off (black trace).

THE ISSUE OF BUNCH LENGTH

The LHC design called for a 1 ns bunch length at the start of physics. In 2011 we have operated with 1.2-1.25 ns, and have observed unexpected heating of some machine components (beam screen, kicker septa, collimator jaws), with a correlation with bunch length. For 2012 the best compromise must be found between heating and luminosity.

The bunch length affects luminosity via the geometric factor

$$F = \frac{1}{\sqrt{1 + \left(\frac{\theta \sigma_z}{2\sigma^*}\right)^2}} \quad (1)$$

Table 1 gives a selection of parameters at the beginning of physics and the resulting geometric factor.

Year	Design length	2011	2012			2016	
Energy (TeV)	3.5	3.5	4	4	4	7	7
β^* (cm)	100	100	60	60	60	55	55
Normalized Transverse Emittance (μm)	2.5	2.5	2.5	2.5	2.5	2.5	2.5
Transverse rms beam size σ^* (μm)	25.88	25.88	18.75	18.75	18.75	13.6	13.6
Full crossing angle θ (μrad)	260	260	300	300	300	250	250
4σ bunch length (ns)	1	1.25	1.25	1.35	1.4	1.25	1.4
Geometric factor F	0.94	0.90	0.80	0.78	0.77	0.76	0.72

Table 1: Geometric factor F for a selection of parameters at the beginning of physics.

With the 1 ns design value (first column) we would have lost 6% luminosity from the crossing angle. In 2011 (second column), we have lost 10% of luminosity due to the 1.25 ns bunch length at the start of physics. At 4 TeV in 2012, with a 60 cm β^* we would lose 20% of luminosity (1.25 ns bunch length), and 23% with a 1.4 ns length.

With longer bunches we have particles close to the separatrix (limit of the RF bucket) and can expect worse longitudinal lifetime. An MD session was conducted in 2011, with eight bunches per beam, identical bunch intensity but different length, circulating at 3.5 TeV without interaction [4]. The beams were separated at the intersection points. Figure 5 shows the loss rate, as a function of the bunch length. The scatter is large but there is a clear correlation of loss with length. In that experiment, we could circulate 1.6 ns long bunches with losses at 1.2 %/h, twice higher than 1.2 ns (0.5 %/h).

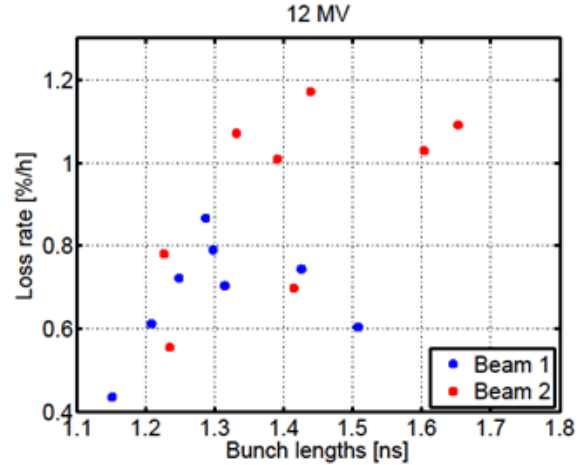


Figure 5: Loss rate vs. bunch length measured during an MD session with 8 bunches per beam at 3.5 TeV, same intensity but different lengths. Beams separated.

In physics we have never observed bunches growing beyond 1.40 ns. The interactions apparently modify the picture: figure 6 shows the lengthening for four different bunches during a physics fill. The bunches were carefully selected to have zero, two, three or four interaction points (IP) per turn (labeled by the indices of the corresponding IPs on figure 6). While the bunch with zero interaction grows to 1.4 ns (black trace) in the fifteen hours of physics, the number of IPs seems to reduce the lengthening, and the bunch colliding at all four IPs – 1528- grows to 1.25 ns only (grey trace).




Figure 6: Bunch lengthening for four bunches in physics. Bunches were chosen to have zero, two, three or four interactions per turn. The bump after seven hours is caused by a cavity tripping and restarting. Reproduced from [5].

Our belief is that the interactions at the IPs lead to a saturation in bunch lengthening long before the RF bucket is full. Figure 7 shows the normal debunching during a fifteen hours long physics fill. The population in the $3\mu\text{s}$ long abort gap is monitored. It gives an indirect measurement of the debunching. If not supplied with new particles, the abort gap gets cleaned in fifteen minutes, time needed for the debunched particles to loose enough energy through synchrotron radiation and move to the momentum collimator. After five hours in physics, the abort gap population has reached an equilibrium between

debunching and momentum collimation. This happens when the bunch length reaches 1.35 – 1.4 ns. Thereafter we observe a clear saturation in bunch lengthening. These observations suggest that the effective momentum aperture is reduced by the beam-beam effect, compared to the bucket half height: bunch lengthening is replaced by particle loss. With a 1.4 ns long bunch, this regime will be reached from the start of physics.



Figure 7: Oct 16-17, 2011. A long fill: more than 15 hours in collision. Beam current (Fast BCT) in green and beige, bunch length mean (red and orange) and abort gap population (blue and violet).

LONGITUDINAL BLOW UP

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 [6]. Longitudinal blow-up provides the required emittance growth [7]. We inject band-limited RF phase-noise in the main accelerating cavities during the whole ramp of about eleven minutes. 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. Only with this feedback, could we reproducibly achieve the programmed bunch length.

As most hadron machines, the LHC Low Level RF (LLRF) includes a main phase loop that measures the beam to accelerating field phase error, averaged over all bunches at every turn, and corrects the RF to minimize the difference. In a hadron collider, this loop is essential to limit the diffusion caused by the RF phase noise at, and around, the synchrotron frequency f_{s0} [8]. While intended to cancel the effect of unwanted external noise sources, this loop would also null the effect of the phase noise injected into the cavity at $f_{RF} \pm f_{s0}$ for the controlled blow-up. This noise signal was therefore injected as an offset in the main phase loop, hopefully resulting in the desired spectrum in the cavity field. As the main phase loop measures the bunch behaviour, it is most sensitive to the behaviour of the populated core of the bunch, and the feedback would therefore much reduce the excitation in the corresponding frequency range. The less populated tails would return a weaker signal, therefore leaving the excitation at full level. The result is a much distorted spectrum, with small excitation at the core and violent excitation at the edges of the bunch (figure 8).




Figure 8: Measured Power Spectral Density of the RF phase noise, during controlled blow-up, towards the end of the acceleration ramp. The synchrotron frequency at the core of the bunch is 28 Hz. The intended excitation spectrum is flat from 24 Hz to 30.8 Hz. The spectrum measured in the cavity is not! The distortion comes from the response of the main phase loop that reduces the excitation in frequencies corresponding to the most populated parts of the bunch.

As we average the phase error over all bunches, the main phase loop only damps the dipole oscillation mode zero[‡]. It will not interact with the excitation of the higher dipole modes, that is a noise spectrum around $f_{RF} \pm n f_{rev} \pm f_{s0}$ (with $n \neq 0$). After a first test in 2010, this scheme was refined during the ions run on December 7th, 2011 with excitation on the first revolution frequency band ($n=1$). At the time we had 358 Lead ion bunches per ring, of low intensity, so that stability is preserved without blow-up in the ramp. Figure 9a shows the length evolution bunch per bunch.




Figure 9a: Blow-up on the first revolution frequency side-band (from time 02:10 till 02:15) with 358 Lead ion bunches. December 7th, 2011.

The machine gets filled in about 45 minutes, consisting of fifteen successive injections in each ring. The large spread in bunch length is the result of Intra Beam Scattering, very violent for ions at injection energy. Around 01:42 starts the acceleration ramp without blow-up, resulting in a bunch length shrinking down to

[‡] This is strictly true for a bunch pattern symmetric around the ring only. In physics the LHC bunch pattern is close to symmetric.

0.6-1.0 ns[§]. Around 02:10 controlled blow-up is applied at $f_{RF} \pm f_{rev} \pm f_{s0}$ resulting in the regular and very uniform observed bunch lengthening in the five minutes long excitation period. Figure 9b shows the length averaged over all bunches during the five minutes long blow-up. The lengthening is smooth and regular.

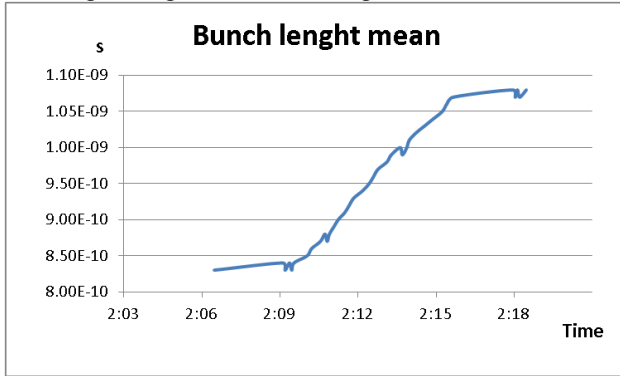


Figure 9b: Bunch length averaged over the 358 bunches during the blow-up.

In the above trial we have kept the excitation amplitude constant, using no feedback from the measured bunch length. The new blow up appears much more predictable than the old one acting through the main phase loop. It opens the road to selective blow-up of the injected batch to reduce IBS driven transverse emittance growth of the idling bunches, during the thirty minutes long filling in 2012 (twelve batch injections per ring).

LONGITUDINAL DAMPER

As mentioned above, the LHC relies on Landau damping for longitudinal stability and no longitudinal damper was planned in its design. Such a system, with a limited bandwidth, would be welcome to damp injection errors though. Indeed, long lasting dipole oscillations have been observed in batch mode (50 ns bunch spacing): growing for ten minutes after injection, then decaying with more than thirty minutes time constant [9] (figure 10).

In 2012, we will use the accelerating cavities as longitudinal kickers, changing the RF phase in the 1 μ s long gap between the circulating beam and the freshly injected batch. The available kick strength (50 kV/cavity) is sufficient to damp the common mode, that is the average phase/energy error of the new batch, before filamentation takes place [10]. The longitudinal damper acting by modulation of the RF field phase looks promising for damping batch-per-batch injection errors but it does not have sufficient bandwidth to act on the bunch-per-bunch phase error in a given batch.

[§] The 1.1 10^{11} p bunch would be unstable with this length, but the Lead ions bunch has a much reduced bunch charge (10^{10} p equivalent) that makes it stable.

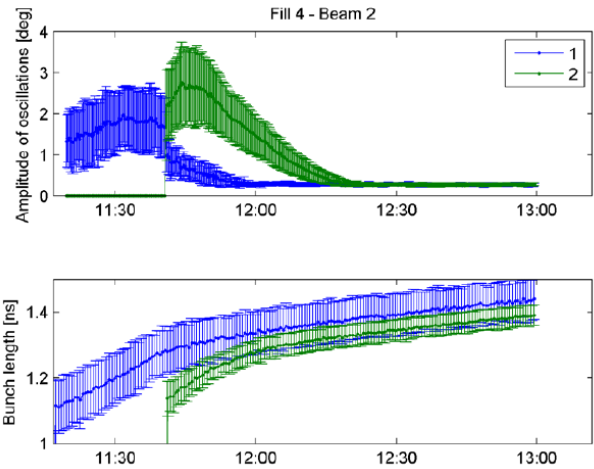


Figure 10: Amplitude of the dipole (top) and quadrupole (bottom) oscillations with batch injection (12 bunches batch followed by a 36 bunch-batch). Reproduced from [9].

THE ROAD TOWARDS MORE LUMINOSITY

Higher energy

Thanks to the longitudinal blow-up that keeps bunch length constant, the stability is actually improved during the acceleration ramp as the voltage rises. At constant bunch length and voltage, it is independent of energy [7][11]. We therefore expect no problem from operation at 4 TeV in 2012.

Higher bunch intensity

In 2011 we have circulated single bunches of intensity $2.85 \cdot 10^{11}$ p at 450 GeV [12]. With the longitudinal blow-up, the stability margin actually improves in the ramp, and we therefore expect no problem related to the broadband longitudinal impedance.

25 ns bunch spacing

In 2011 all physics fills were with 50 ns bunch spacing, but during dedicated Machine Development sessions, we have circulated 2100 bunches at 25 ns spacing, 10^{11} p/bunch, at injection energy [13]. The Low Level RF was designed to reduce the cavity impedance at the fundamental and our calculations indicate that the Landau damping provided by the 4 Hz synchrotron tune spread is sufficient to preserve stability at ultimate beam current (2808 bunches, $1.7 \cdot 10^{11}$ p/bunch) [10]. The damping of the cavity HOMs is also designed for above ultimate. These conclusions rest on calculations only and must be confirmed by measurements of the damping rate of coupled bunch modes during Machine Development sessions in 2012.

The transient beam loading caused by the gaps in beam current is compensated by the Low Level RF at the expense of klystron power. With future nominal beam intensity (2808 bunches, $1.1 \cdot 10^{11}$ p/bunch, 1.2 ns 4σ length), we would need 200 kW in physics

(1.5 MV/cavity, $Q_L=60000$) and 110 kW during filling (0.75 MV/cavity, $Q_L=20000$). With the present DC settings (50 kV, 8A), the klystrons saturate at 200 kW RF. For nominal beam operation we plan to change the DC settings, at the end of filling and before starting the ramp. With 57 kV and 8.7 A, the klystrons saturate at 270 kW RF, providing sufficient margin for regulation with respect to the 200 kW threshold. The change of klystron DC settings with beam circulating at injection energy, followed by the ramp, was tested successfully with 60 bunches per ring at 25 ns spacing in 2011[13]. It must be validated with longer batches in 2012.

With ultimate conditions (2808 bunches, 25 ns spacing, $1.7 \cdot 10^{11}$ p/bunch) planned after LS1, the present scheme requires more than 300 kW in physics (12 MV), that are not available from the klystron. We would then accept the modulation of the cavity phase by the beam current (transient beam loading) and adapt the set point for each bunch accordingly. This mode of operation was planned during the design of the LHC RF [14]. Stability will be preserved and we would need 105 kW only, (in theory) independent of the beam current. The penalty is a modulation of the cavity phase that changes the bunch spacing and therefore the collision point. Figure 11 shows the resulting phase slip along the ring: the 65 ps displacement is small compared to the 1.2 ns 4σ bunch length. As the filling pattern of the two rings is very similar, the phase modulations will cancel out in IP1 and IP5 and the resulting displacement of the collision vertex will be much smaller than the above 65 ps. This scheme will be tested during Machine Development sessions in 2012 to prepare the way for after LS1.




Figure 11: Modulation of the cavity phase by the transient beam loading in physics. 2835 bunches, $1.7 \cdot 10^{11}$ p/bunch, 1.5 MV/cavity, $Q_L=60k$. The abort gap spans 127 empty buckets (25 ns spacing) or $3.2 \mu s$. Filling as in the original LHC design.

CONCLUSIONS

In 2012 the LHC RF intends the following upgrades

- Further reduction of capture losses with the longitudinal damper.
- Modification of the blow-up method.
- Batch per batch blow-up at injection.

Data indicate that in physics the momentum aperture is smaller than the bucket half height. With 1.4 ns, the effective bucket is full. If (and only if) it is really needed for other equipment, we propose to test physics with slightly longer bunches, early in the re-start.

Thanks to the longitudinal blow-up, the stability is independent of the energy. Operation at 4 TeV should not cause problems.

We have not approached the single bunch intensity limit. We have circulated $3 \cdot 10^{11}$ p/bunch with no observed instabilities.

There was no surprise from the first 25 ns tests.

The RF can deal with nominal total intensity (2808 bunches, 25 ns, $1.1 \cdot 10^{11}$ p/bunch)

- On the stability side, calculations show large margin from HOMs and cavity impedance at the fundamental. However we would like to measure damping time of coupled-bunch mode for confirmation. MD time needed.
- On the klystron power side, we can deal with nominal if we increase the klystron DC settings (High Voltage) before the ramp (to be tested during 25 ns MDs).

For ultimate intensity (2808 bunches, 25 ns, $1.7 \cdot 10^{11}$ p/bunch), the RF must allow for the modulation of the cavity phase by the transient beam loading. First tests should take place before LS1.

ACKNOWLEDGEMENTS

Many thanks to T. Bohl for the explanations on the SPS longitudinal parameters, and to R. Calaga for the data on geometric factor versus bunch length. We are very grateful to T. Bohl, E. Shaposhnikova and J. Tuckmantel for the many discussions on beam diffusion and longitudinal blow-up. The LHC RF features sophisticated electronics (longitudinal blow-up, one-turn feedback, longitudinal damper) that are made possible by the competence and dedication of our colleagues J. Molendijk (on the firmware side) and M. Jaussi (for the controls issues). We count on them all for even more success in 2012.

REFERENCES

- [1] P. Baudrenghien et al., The LHC RF: 2011 and beyond, Evian 2011
- [2] T. Bohl et al, Controlled Longitudinal Emittance Blow-up in the SPS as injector and LHC Test-Bed, CERN-AB-Note-2003-084-MD, Dec 2003
- [3] T. Mastoridis et al., The LHC One-Turn Feedback, CERN-ATS-Note-2012-025 PERF
- [4] J.F. Esteban Muller et al., Dependence of single beam lifetime on bunch length, ATS-Note-2011-083 MD
- [5] G. Papotti, Colliding beams: luminosity and beam-beam, LHC Beam Operation workshop, Evian 2011
- [6] E. Shaposhnikova et al., Loss of Landau damping in the LHC, IPAC 2011
- [7] P. Baudrenghien et al., Longitudinal Emittance Blow-up in the LHC, IPAC 2011

- [8] T. Mastoridis et al., Radio frequency noise effects on the CERN Large Hadron Collider beam diffusion, Phys. Rev. Spec. Topics Accel. Beams.,14, 092802 (2011)
- [9] P. Baudrenghien et al., Longitudinal Oscillations with Batch Injection in the LHC, ATS Note 2011/031/MD
- [10] P. Baudrenghien et al., The LHC RF: Is it working well enough? Chamonix 2011 workshop
- [11] E. Shaposhnikova, longitudinal beam parameters during acceleration in the LHC, LHC Project Note 242, Dec 2000
- [12] T. Mastoridis et al., RF observations during High Single Bunch Intensity MD, CERN-ATS-Note-2012-023 MD
- [13] T. Mastoridis et al., RF observations during 25 ns MD sessions in 2011, CERN-ATS-Note-2012-024 MD
- [14] D. Boussard, RF Power Requirements for a High Intensity Proton Collider, PAC 1991