STUDIES ON CONTROLLED RF NOISE FOR THE LHC

H. Timko*, P. Baudrenghien, E. Shaposhnikova, CERN, Geneva, Switzerland
T. Mastoridis, California Polytechnic State University, San Luis Obispo, USA

Abstract

RF phase noise is purposely injected into the LHC 400 MHz RF system during the acceleration ramp for controlled longitudinal emittance blow-up, in order to maintain longitudinal beam stability. Although the operational blow-up works reliably, studies of the injected RF noise are desirable not only to allow for a better-controlled, more flexible blow-up, but also for other applications such as the mitigation of machine-component heating through appropriate bunch shaping. Concerning the noise injection, an alternative algorithm was developed and implemented in the hardware, but first tests revealed unexpected modulation of the achieved bunch length along the ring, and subsequently, theoretical studies have been launched. In this paper, we present a summary of ongoing measurement analysis and simulation studies that shall explain previous observations, predict what can be expected in different cases, and thus help to optimise the RF noise in general.

INTRODUCTION

Controlled longitudinal emittance blow-up is necessary in the LHC to maintain longitudinal beam stability during the acceleration ramp. For a constant stability margin, the emittance ε should be increased with energy as $\varepsilon \propto \sqrt{E}$ [1], resulting in a roughly constant bucket filling factor and relative synchrotron frequency spread. In practice, a good blow-up was achieved by a feedback that scales the RF phase noise $\varphi_N(t)$ to keep the bunch length constant [2, 3]. At flat top, also bunch shaping through noise injection can be desirable, for instance to mitigate machine-component heating.

The noise spectrum $S_\varphi(f)$ applied determines in what frequency range diffusion is triggered in the bunch, and hence, determines also the resulting bunch shape. Diffusion under external noise [4–6] and bunch shaping with band-limited white noise [7] has been studied in the past. Controlled emittance blow-up for the SPS and LHC was designed and implemented [2, 8] subsequently. In the LHC, a band-limited white noise spectrum is applied, however, feedback loops complicate the analysis. A constant relative noise band is chosen with the range $(0.86–1.1) f_{s0}$ (Fig. 1a), where f_{s0} is the synchrotron frequency of the synchronous particle. In a single-RF system this will affect the core of the bunch that has a maximum phase amplitude of synchrotron oscillations in the range of $\varphi_{\text{max}} = (0–1.52)$ rad (Fig. 1b).

Operationally, $\varphi_N(t)$ is applied through the phase loop (PL), as an additional phase shift to the phase correction between bunch and RF phase [2]. This works reliably, and even for a full machine the beam can be brought stably through the acceleration ramp, with a typical target bunch length of 1.2 ns achieved homogeneously (±30 ps) for all bunches [2]. However, since the phase loop corrects the centre-of-mass motion, the effective noise spectrum is reduced greatly around f_{s0}, see Fig. 1a. This modification of the spectrum makes it difficult to shape bunches in a well-determined way.

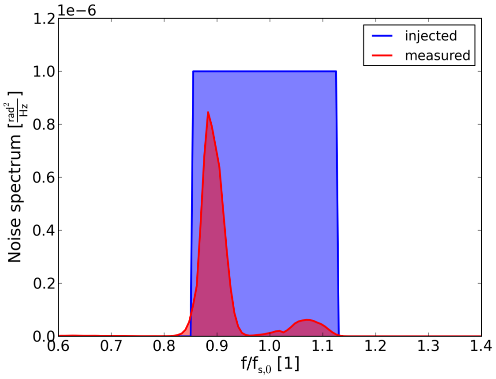

(a) Injected (blue) and measured (red) noise power spectral density. The measured noise is applied through the PL.

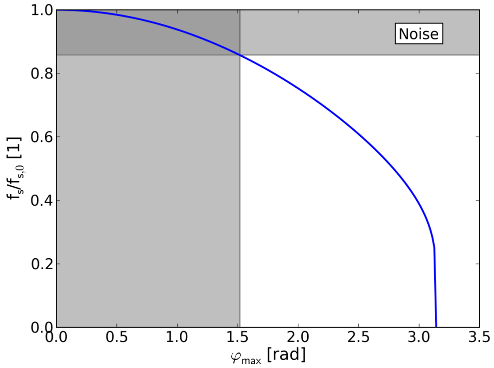

(b) Synchrotron frequency distribution in a single-RF system as a function of maximum phase coordinate. The dashed region marks the range affected by the injected noise.

Figure 1: LHC controlled longitudinal emittance blow-up.

Alternatively, $\varphi_N(t)$ can be applied through the cavity controller (CC). The PL is still required for the ramp, but an interaction between the noise and the PL can be avoided: using a symmetric filling pattern, and injecting the noise on a revolution frequency side-band, the noise is practically invisible to the PL. This scheme was tested in early 2012, and indeed, $S_\varphi(f)$ in the cavity field reproduced the desired spectrum exactly [9]. At the time, however, this noise injection scheme could not be made operational, because tests with few bunches led to a bifurcation of the final bunch length: although on average the target bunch length was obtained, the bunches were either too short or too long. Further studies are planned after the start-up of the LHC early 2015.

---

* helga.timko@cern.ch

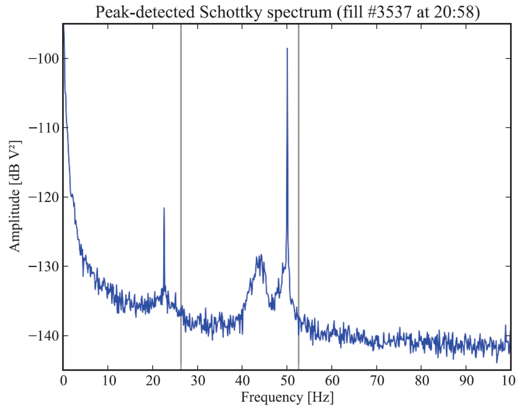

Figure 2: Schottky spectrum of LHC bunches at flat top energy (from [10], measured at 4 TeV, 12 MV). The vertical lines mark f_{s0} = 26.3 Hz and $2f_{s0}$ = 52.6 Hz.

Another motivation for our studies is that the peak-detected Schottky signal of bunches at flat top revealed a 'hole' in the bunch distribution somewhat below $2f_{s0}$ (see Fig. 2, [10]). Note that, using this measurement technique, the dipole band is distorted and the quadrupole band gives the best reproduction of the particle distribution [11]. In principle, it is not excluded that the 50 Hz line could be related to creating this hole, as the core gets depopulated when the 50 Hz line is crossed [12]. However, this happens quite early during the acceleration ramp (f_{s0} = 55.2 Hz and 24.0 Hz at flat bottom and flat top, respectively), and during the ramp the emittance is blown up roughly by a factor 4, using a noise spectrum that follows f_{s0} and thus mixes different frequencies. Both these facts make it unlikely that the observed depopulation is due to the 50 Hz line. Instead, it could be related to the operational emittance blow-up applied through the PL. Understanding the origin of this is important because bunches with such a distribution are intrinsically more unstable and result also in a different heating of machine components.

SIMULATION MODEL

The CERN BLonD simulation code [13] has been extended to contain a complete model of the LHC controlled noise injection during the acceleration ramp. For the time being, the code is restricted to single-bunch simulations. The incoherent synchrotron frequency shift in the LHC ($\sim 0.01 f_{s0}$ [2]) is negligible compared to the noise bandwidth, and thus, intensity effects were neglected.

The phase noise $\varphi_N(t)$ was generated in a similar fashion as in the LHC [14]. The double-sided power spectral density $S_\varphi(f)$, $[S_\varphi(f)] = \frac{\mathrm{rad}^2}{\mathrm{Hz}}$, is transformed to $\varphi_N(t)$ via a real FFT, such that one data point per turn is obtained. Just like in measurements, a flat spectrum between $(0.86–1.1) f_{s0}$ is applied. Since f_{s0} and the revolution period change during the ramp (Fig. 3), the spectrum is adjusted every 10,000 turns and $\varphi_N(t)$ is re-generated with a new seed for the random number generator. Furthermore, the amplitude A_S of the spectrum $S_\varphi(f)$ is scaled such that the r.m.s. phase noise $\varphi_N^{\mathrm{rms}} = \sqrt{\int S_\varphi(f) df} = \sqrt{2 \times 0.24 f_{s0} A_S}$ remains constant. We compare two cases, $\varphi_N^{\mathrm{rms}} = 2°$ and $0.2°$, with the latter being more realistic for the ramp. $\varphi_N^{\mathrm{rms}} = 0.2°$ results in $A_S = 4.4 \times 10^{-7}$ rad^2/Hz and $A_S = 1.0 \times 10^{-6}$ rad^2/Hz at flat bottom and flat top energies, respectively.

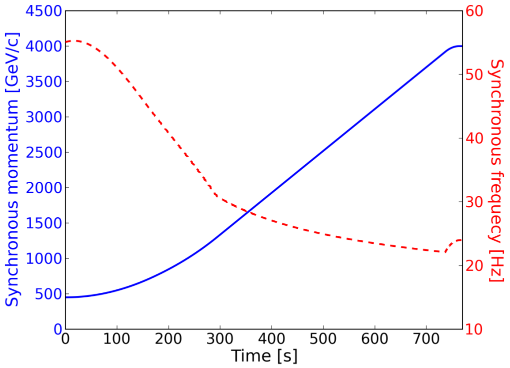

Figure 3: Beam momentum and synchrotron frequency during the LHC acceleration ramp, used both in the machine and in simulations. The voltage is increased linearly from 6 MV to 10 MV during this time.

Also the PL is modelled in simulations according to the real LHC implementation [15]. The PL calculates every turn the phase difference $\Delta\varphi_{\mathrm{PL}} = \varphi_{\mathrm{COM}} - \varphi_s$ between the bunch centre-of-mass φ_{COM} and the RF phase, which in this case is substituted by the synchronous phase φ_s calculated from the design beam momentum and RF voltage. In the equations of motion, the RF angular frequency is then corrected in the subsequent turn by the PL to $\omega_{\mathrm{RF}}^{(n+1)} = h\omega_{s0}^{(n+1)} - g\Delta\varphi_{\mathrm{PL}}^{(n)}$, where $h\omega_{s0}$ is the design RF angular frequency and $g = 1/(5\ \mathrm{turns})$ is the PL gain. $\varphi_N(t)$ is then added either to the phase correction φ_{PL} (PL case) or directly in the energy kick in the equations of motion (CC case).

During acceleration, an additional feedback on the bunch length is acting, which regulates the phase noise amplitude to maintain a constant bunch length. The feedback 'measures' the bunch length $\tau_{\mathrm{meas}} = 4/2.355\tau_{\mathrm{FWHM}}$ [1] every 3 s (33,740 turns) and compares it to the target 4-sigma bunch length τ_{targ} = 1.2 ns. For the subsequent 3 s, $\varphi_N(t)$ is multiplied by the factor x determined through the recursion [2] $x^{(n+1)} = ax^{(n)} + k(\tau_{\mathrm{targ}} - \tau_{\mathrm{meas}})$, with $x^{(0)} = 1$, limiting x to the range $[0, 1]$, and using the same constant $a = 0.8$ and gain 1.5×10^9 s^{-1} as in the LHC.

SIMULATION RESULTS

Flat Bottom Simulations

First, we describe simulation results of noise injected at flat bottom energy. These simulations were done without the bunch-length feedback. Noise was injected during 88.9 s (10^6 turns) and the bunch was relaxed during 8.89 s (10^5 turns) subsequently. Both injection through PL and CC, as well as $\varphi_N^{\mathrm{rms}} = 0.2°$ and 2° were studied. In addition, to

---

[1] For a Gaussian bunch, FWHM = $2\sqrt{2 \ln 2}\sigma \approx 2.355\sigma$. Hence, we scale by 4/2.355 to obtain the corresponding 4-sigma bunch length.

Table 1: Comparison of r.m.s., Gaussian fit, and FWHM Bunch Lengths for Different Blow-up Settings at Flat Bottom. All bunch length values are 4-sigma. The FWHM value is scaled such that all three methods give the same result for a Gaussian bunch. The time needed to reach steady-state is shown as well. The RF bucket length is 2.5 ns.

Loop	φ_N^{rms}	**Seed**	4× **r.m.s.**	**4-sigma Gaussian fit**	4/2.355× **FWHM**	**Time constant**
Initial distribution			1.26 ns	1.22 ns	1.09 ns	–
PL	0.2°	multi-seed	1.27 ns	1.26 ns	1.18 ns	> 100 ns
		single-seed	1.27 ns	1.25 ns	1.18 ns	85 s
	2°	multi-seed	1.46 ns	1.68 ns	1.97 ns	3 s
		single-seed	1.41 ns	1.62 ns	1.85 ns	4 s
CC	0.2°	multi-seed	1.50 ns	1.72 ns	1.97 ns	25 s
		single-seed	1.41 ns	1.60 ns	1.76 ns	15 s
	2°	multi-seed	1.75 ns	2.03 ns	2.27 ns	6 s
		single-seed	1.53 ns	1.78 ns	2.01 ns	1.2 s

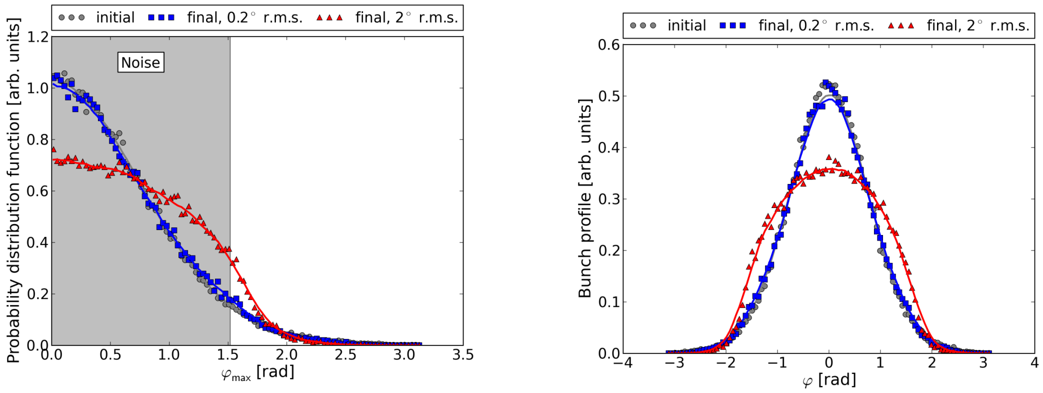

(a) Noise injection through the PL with 0.2° (blue) and 2° (red) r.m.s. phase noise. Initial distribution in grey.

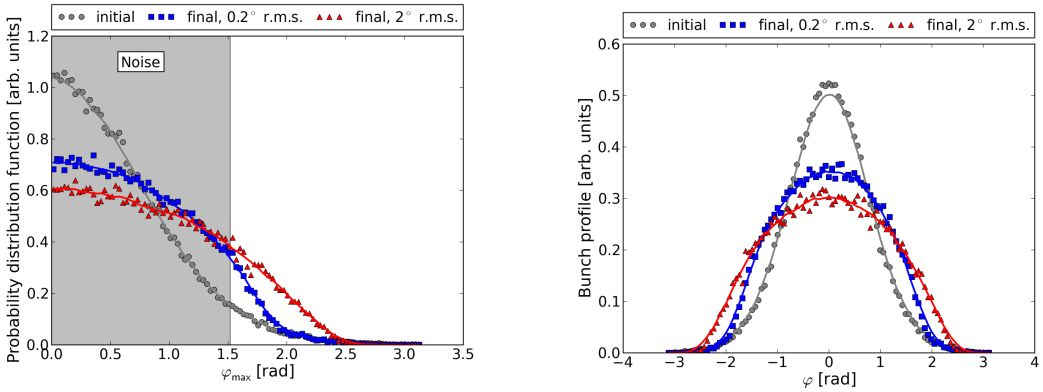

(b) Noise injection through the CC with 0.2° (blue) and 2° (red) r.m.s. phase noise. Initial distribution in grey.

Figure 4: Injection of multi-seed noise at flat bottom. Comparison of initial and final probability density functions (left) and bunch profiles (right).

study the effect of periodicity in the random number generator, the noise was generated from the same distribution in two different manners: once using a single seed, and once changing the seed every 10,000 turns (as done during the ramp). All simulations presented here were using the same initial matched distribution of 50,001 macro-particles.

The initial and final bunch profiles as well as the probability distribution functions as a function of the maximum phase coordinate are shown in Fig. 4. When the noise is injected through the PL, the core population is clearly higher than in the CC case. The difference is even more prominent with $\varphi_N^{\mathrm{rms}} = 0.2°$. This reflects well the fact that the PL reduces the noise spectrum around f_{s0} significantly (cf. Fig. 1a). For the same injection method, the factor 10 difference in φ_N^{rms} has a dramatic impact on the final bunch shape. Comparing the PL, $\varphi_N^{\mathrm{rms}} = 2°$ case with the CC, $\varphi_N^{\mathrm{rms}} = 0.2°$ case, it requires about this factor of 10 in phase (factor 100 in A_S) to compensate for the difference between PL and CC.

In Figure 4, we showed only the multi-seed cases. The final distributions of the single-seed cases are very similar, apart from being blown up less. The main difference between the single- and multi-seed cases is in the blow-up efficiency reflected by the final bunch length, see Tab. 1. This can be explained by the single-seed case introducing numerically a periodicity in the phase noise, which creates ‘islands’ in the phase space of the bunch. With multiple seeds, on the other hand, one can ensure that the complete phase-space region targeted by the noise spectrum is indeed affected. Hence, the bunch is blown up more.

Given that the bunch distribution is changing during the blow-up, it is non-trivial to choose a good measure of the bunch length. In Tab. 1 we compare three different methods: r.m.s., Gaussian fit, and FWHM. While r.m.s. is the natural measure to track bunch lengthening due to diffusion, with the Gaussian fit or FWHM values additional information on the bunch shape can be obtained. Comparing r.m.s. and Gaussian fit values, for instance, one can see that the initial distribution is close to Gaussian, but the final distribution is far from it. The most sensitive measure is FWHM, whose value is the lowest for the initial and the highest for the final distribution. Indeed, also in the LHC it is chosen for the bunch-length feedback because it reflects well the average ‘curvature’ (shape) of the bunch.

When applying a noise of constant spectrum, the bunch will diffuse during a given time span, after which the bunch profile reaches a steady state, even though the bunch is still being shaken. The typical time evolution of the (r.m.s.) bunch length is shown in Fig. 5. The bunch length saturates with a $ae^{-t/\tau}+b$ trend; the time constants τ are summarised in Tab. 1 as well. A theoretical estimate of the time constants to expect in the different cases is not easy to obtain, because (i) the noise spectrum is band-limited, (ii) the PL distorts the spectrum, and (iii) the short-bunch approximation cannot be applied either. However, we can draw a few important conclusions from the data:

- the noise realisation (seeding) can affect the bunch length growth rate significantly,
- the multi-seed realisation blows up the bunch more and with a longer time constant,
- for $\varphi_N^{\mathrm{rms}} = 2°$, the time constants are a few hundred synchrotron periods only and become comparable for the PL/CC cases,
- $\varphi_N^{\mathrm{rms}} = 0.2°$, the difference between PL and CC both in final bunch length and time constant is significant.

Simulations During the Acceleration Ramp

In the following, we present some first results of simulations with the acceleration ramp, applying $\varphi_N^{\mathrm{rms}} = 0.2°$ either through the PL or the CC. In both cases, the bunch-length feedback was on; the evolution of the phase-noise scaling factor and the resulting bunch length during the ramp are shown in Fig. 6. When injecting the noise through the CC, $\tau_{\mathrm{targ}} = 1.2$ ns can on average be met throughout the whole ramp, with about (20–40) % of $\varphi_N(t)$. Note that the 4-sigma r.m.s. bunch length deviates from τ_{targ} as the shape of the bunch is changing. However, $\varphi_N^{\mathrm{rms}} = 0.2°$ is insufficient to blow up the beam through the PL; even with 100 % of $\varphi_N(t)$, the bunch length shrinks below 0.7 ns. Qualitatively, this is in line with the results from flat bottom simulations. Further simulations with increased φ_N^{rms} are underway.

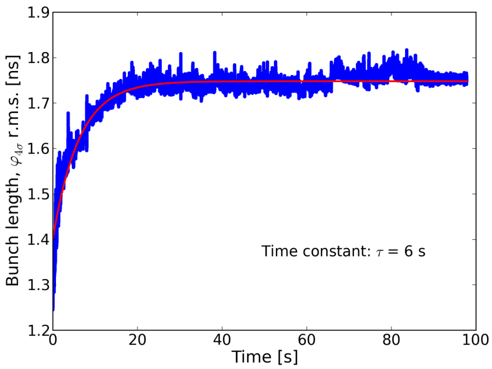

Figure 5: Typical evolution of 4-sigma r.m.s. bunch length over time (blue) and the fit of the form $ae^{-t/\tau}+b$ (red).

The 4-sigma bunch lengths at the end of the ramp are: 0.66 ns r.m.s., 0.65 ns Gaussian fit, and 0.63 ns full-width ($4/2.355\times$ FWHM) with the PL; 1.05 ns r.m.s., 1.15 ns Gaussian fit, and 1.18 ns full-width with the CC. As expected, the shrinking bunch with insufficient blow-up remains Gaussian. On the other hand, the bunch that is blown up through the CC is rounder, with a denser core up to ±1.2 rad, sharply decreasing till ±1.7 rad, with no tail population at all. The low tail population is consistent with earlier observations [9] and is due to the mixture of controlled emittance increase and decreasing bucket filling factor during the ramp, cf. Fig. 7. The initial and end-of-ramp (8.5×10^6 turns, 756 s) bunch profiles are shown as well.

Looking closer at the right-hand-side plot in Fig. 6, one can see that the bunch length follows the same pattern as the scaling factor x of the bunch-length feedback. Also, $\varphi_N(t)$ needs to be adjusted much more in the beginning of the ramp than later.

CONCLUSIONS AND OUTLOOK

Noise generation, phase loop, and bunch-length feedback have been implemented in the CERN BLonD code, creating a valuable tool for the realistic simulation of controlled bunch shaping and emittance blow-up during the ramp in the LHC. First simulation results at flat bottom and during the acceleration ramp have been presented. Simulations reproduce the reduced efficiency of the phase-loop-injected noise in the core of the bunch.

The blow-up results in a round bunch profile that is far from Gaussian. Bunch lengths obtained with different methods – r.m.s., Gaussian fit, and FWHM – differ largely. Hence, also for on-line monitoring of the bunch length in the LHC, a suitable choice of bunch length definition is important. Out of the three studied, the FWHM measure is the most sensitive to the bunch shape.

First simulations of noise injection during the acceleration ramp confirm the experimental observations that after the blow-up the bunch has a rounder core and lower tail

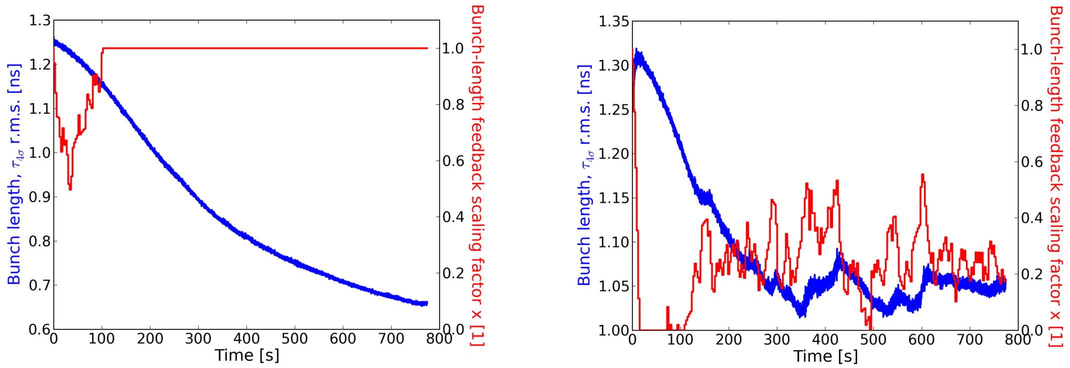

Figure 6: Bunch-length feedback scaling factor x and resulting 4-sigma r.m.s. bunch length during the ramp, with PL (left) and CC (right).

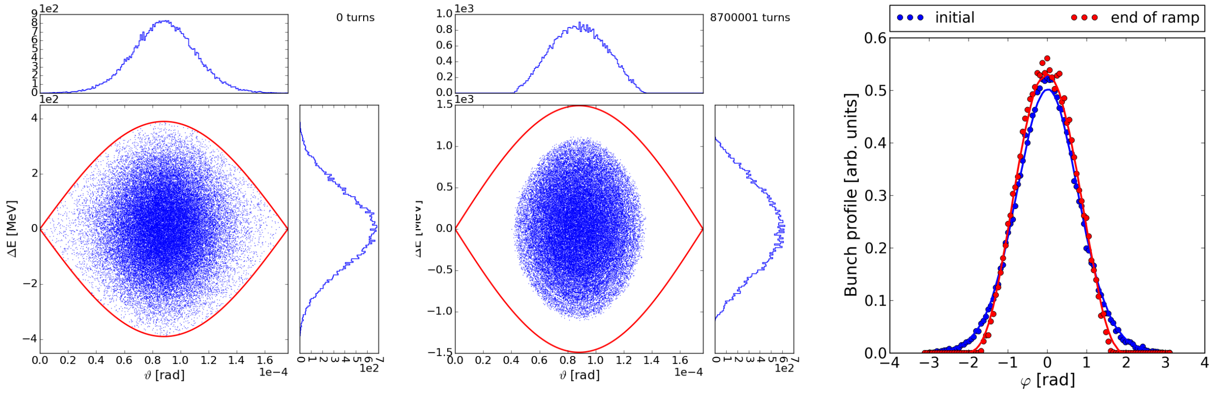

Figure 7: Initial (left) and final (middle) phase space and bunch profiles (right); $\vartheta = \varphi/35640$. Blow-up through the CC.

population. To obtain the same final bunch length, different noise amplitudes for injection through phase loop and cavity controller are required, since the phase loop reduces the effective noise spectrum.

Studies of controlled noise injection during the ramp will be continued to optimise the blow-up for the next LHC run. To better model the noise injection through the cavity controller, multi-bunch simulations are planned in the long term. This requires not only the implementation of multi-bunch capabilities in BLonD, but also an overall optimisation of the run-times for demanding simulations such as the LHC ramp. The LHC re-commissioning early next year will give a good opportunity to test new schemes and settings in real life.

REFERENCES

[1] E. Shaposhnikova, “Longitudinal beam parameters during acceleration in the LHC”, LHC Project Note 242, Geneva, (2000).

[2] P. Baudrenghien et al., “Longitudinal emittance blow-up in the LHC”, IPAC‘11, San Sebastian, Spain, TUPZ010 (2011).

[3] E. Shaposhnikova, et al., “Loss of Landau damping in the LHC”, IPAC‘11, San Sebastian, Spain, MOPC057 (2011).

[4] S. Krinsky, J.M. Wang, “Bunch diffusion due to RF noise”, Part. Accel. 12, 107–117 (1982).

[5] G. Dôme, “Diffusion due to RF noise”, CERN Accelerator School ‘85, Oxford, UK, 370–401 (1985).

[6] S.V. Ivanov, “Longitudinal diffusion of a proton bunch under external noise”, IHEP Preprint 92-43, Protvino (1992).

[7] T. Toyama, “Uniform bunch formation by RF voltage modulation with a band-limited white signal”, NIM A, 447, 317–327 (2000).

[8] J. Tückmantel et al., “Experimental studies of controlled longitudinal emittance blow-up in the SPS as LHC injector and LHC test-bed”, LHC Project Report 761, Geneva, (2004).

[9] P. Baudrenghien, T. Mastoridis, “Longitudinal emittance blowup in the large hadron collider”, NIM A, 726, 181–190 (2013).

[10] T.E. Levens et al., “A beam-synchronous gated peak-detector for the LHC beam observation system”, IBIC‘13, Oxford, UK, MOPC35 (2013).

[11] E. Shaposhnikova, et al., “Longitudinal peak detected Schottky spectrum”, HB‘10, Morschach, Switzerland, TUO1C04 (2010).

[12] J. Tückmantel, “Simulation of LHC bunches under influence of 50-Hz-multiple lines on the cavity field”, LHC Project Note 404, Geneva, (2007).

[13] BLonD Beam Longitudinal Dynamics simulation code website: `http://blond.web.cern.ch`

[14] J. Tückmantel, “Digital generation of noise-signals with arbitrary constant or time-varying spectra”, LHC Project Report 1055, Geneva (2008).

[15] P. Baudrenghien et al., ‘The LHC low level RF’, EPAC‘06, Edinburgh, UK, TUPCH195 (2006).