

ELECTRON CLOUD EFFECTS IN THE LHC IN 2011

G. Rumolo*, G. Iadarola, O. Domínguez, G. Arduini
H. Bartosik, S. Claudet, J. Esteban-Müller, F. Roncarolo, E. Shaposhnikova, L. Tavian

Abstract

The LHC began operation with 50ns beams in early April, 2011. The observation of pressure rise, heat load in the arcs, beam quality degradation and synchronous phase shift during the first days with 50ns beams clearly revealed the development of an electron cloud inside the machine. However, a 5-day dedicated scrubbing run with 50ns beams was sufficient to mitigate this effect and allow for subsequent physics production with this bunch spacing. Furthermore, from the end of June to the end of October, 2011, five Machine Development (MD) sessions with 25ns beams took place, during which all the aforementioned electron cloud indicators again made an appearance and further scrubbing could be achieved.

In this paper, we will first briefly explain how the electron cloud can be detected in the LHC and give a summary of the electron cloud observations in the LHC during 2011. Then, we will quantify the scrubbing process in terms of evolution of secondary electron yield (SEY or δ_{\max}) of the chamber surface by comparing the measured data — pressure, heat load, synchronous phase shift — with the simulation results, obtained with newly developed and optimized tools. Finally, we will discuss the influence of the electron cloud on the evolution of some beam observables.

INTRODUCTION

When an electron cloud builds up in a running accelerator, the beam chamber becomes filled with an electron gas whose distribution and flux to the walls mainly depend on the beam structure and the properties of the beam chamber (i.e. geometry and δ_{\max} of the inner surface) [1]. The flux of electrons hitting the wall of the vacuum chamber and its energy distribution, $\phi_e(E)$, are the origin of both pressure rise, ΔP , and power deposition on the chamber wall, ΔW , which can be expressed as:

$$\Delta P = kT \frac{\int \eta_e(E) \phi_e(E) dE}{S_{\text{eff}}} \quad (1)$$

$$\Delta W = \int \phi_e(E) E dE \quad (2)$$

Here k is the Boltzmann constant, T the temperature, S_{eff} the local pumping speed, η_e the desorption yield, E the electron energy. While the pressure rise is directly measurable through the vacuum gauges, the power deposited by the electrons on the chamber walls can be either revealed

as the heat load on the wall — proportional to its temperature rise ΔT in absence of cooling — or as the energy per turn lost by the beam. Since the RF system compensates for this energy loss, even at flat bottom the bunches will move to an accelerating stable phase. The change in the stable phase, $\Delta\varphi_s$, can be used to quantify the beam power transfer to the electron cloud, and eventually to the wall.

Furthermore, the presence of a non-neutral electron plasma with local volume density around the beam, ρ_e (i.e. central density), can lead to coherent single or coupled bunch instabilities or incoherent emittance growth and slow loss. While running with high chromaticity settings at injection energy can be desirable to suppress the coherent mechanisms due to electron cloud, this does not cure — or might even have the adverse effect to further enhance — the incoherent phenomena and lead to stable, but degraded, beams. Another issue is that high chromaticity settings must be used with caution when ramping the beam energy, because, in absence of sufficient Landau damping, they might render the beam unstable as a result of its interaction with the machine impedance [2]. For these reasons, this configuration can be regarded as acceptable for the purpose of machine scrubbing, but is of no use for physics operation.

In the LHC the electron cloud is detected either by pressure rise, heat load and synchronous phase shift, or through its detrimental effects on the beam. Machine scrubbing, i.e. the lowering of the SEY of the wall's inner surface over time due to electron bombardment, can then be qualified through the gradual weakening of these symptoms. However, it is worth noticing that both power loss and pressure rise data should be used with care to extrapolate information on the electron cloud in the machine and its state of conditioning. Firstly, they depend closely on the beam structure — i.e. bunch-by-bunch intensity and length, bunch spacing, number of gaps, length of the gaps — and comparing them for different beam conditions is not straightforward. Secondly, there is another subtle complication, which can be seen from the above equations. The observables related to power loss only depend on the electron flux to the walls and their values can be directly compared across different days to infer a possible change in the SEY. On the other hand, the pressure rise depends on the electron flux as well as on the desorption yield, η_e . Hence, it is reasonable to assume that a series of subsequent relative measurements can be used to characterize the SEY in the neighbourhood of a vacuum gauge at a given time, but measurements done in different days cannot be directly compared, even in the same beam conditions, because any change observed is the result of both the scrubbing in SEY and η_e .

* Giovanni.Rumolo@cern.ch

HISTORICAL

Operation with 50ns beams: scrubbing run and physics

The 2011 run of the LHC started in March with beams with 75ns bunch spacing. Owing to beam scrubbing from the late 2010 MD sessions with 75 and 50ns beams, the LHC could quickly move to physics with 75ns beams without suffering from major outgassing limitations or beam instabilities. In less than one month, the LHC was already able to successfully accelerate and collide two 75ns beams each made of 200 bunches distributed in batches of 24. The full week 5–12 April, 2011, was devoted to the scrubbing run with 50ns beams, although three days were lost due to issues independent of the scrubbing program and the last day was devoted to tests of accelerating/colliding beams with 50ns spacing (up to 246 bunches). The goal was to prepare the machine to switch to 50ns beams and thus extend the luminosity reach for the 2011 run. Over the scrubbing run, the number of bunches injected into the LHC was gradually ramped up to a maximum of 1020 per beam (in batches of 36). Several stores at injection energy with different numbers of bunches took place. During the first days of scrubbing, pressure rise, heat load in the arcs, coherent beam instabilities as well as emittance growth were observed [3]. Nevertheless, the beams could be kept in the LHC at injection energy thanks to the fact that:

- The heat load could be handled by the cooling capacity of the cryogenic system;
- The pressure rise was tolerable or, where needed, the interlock level on the pressure value was temporarily raised so as not to cause beam dump;
- High chromaticity settings, acceptable at injection energy, were used for damping the coherent instability;
- The incoherent emittance growth and the associated intensity loss, mainly affecting the last bunches of each batch, were sufficiently slow as not to trigger the Beam Loss Monitors (BLMs).

In these conditions, the electron cloud produced by the circulating beams served the purpose of scrubbing the inner wall's surface of the LHC beam chambers. The strategy adopted to optimize the scrubbing process consisted of constantly topping the total beam intensity in the LHC with the injection of more trains, such that the vacuum activity, and therefore the electron cloud, could be kept at a constant level. This was expected to efficiently reduce the SEY of the walls to a value eventually at the limit for a significant electron cloud build up (and below the threshold for beam instability at nominal intensity). After approximately 17 effective hours of beam scrubbing time — corresponding to about 72h of beam time — the pressure improved by an order of magnitude throughout the machine. At the end of the scrubbing run, a residual pressure rise was still observed in some cold-warm transitions and straight sections, while in the arcs both the heating of the beam screen and the pressure increase disappeared. Also, no electron cloud

instability or emittance growth was seen to affect the beams towards the end of the scrubbing run.

The success of the scrubbing run was proved by the subsequent smooth LHC physics operation with 50ns spaced beams. Between mid-April and end-June the number of bunches collided in the LHC was steadily increased up to its maximum value of 1380 per beam, while the intensity per bunch and the transverse emittances remained constant at their nominal values (i.e., 1.15×10^{11} ppb and $2.5 \mu\text{m}$). During this whole period, the surface scrubbing naturally continued, as witnessed by an additional one-order-of-magnitude decrease of the dynamic pressure level around the machine. By the end of June, beams in a full machine and colliding for 20h hardly exhibited any emittance growth during their time in stable beam mode. The switch to 50ns beams with lower transverse emittances ($1.5 \mu\text{m}$) allowed the LHC peak luminosity to easily score an additional 50% increase, while not leading to any critical recrudescence of the electron cloud. This could be expected, given the weak dependence of electron cloud formation on the beam transverse emittances. Finally, to push the peak luminosity further up, the intensity per bunch was adiabatically increased to approximately 1.5×10^{11} ppb over the final few months of the run. Again, no significant return of the electron cloud was observed during this phase [4].

However, even during the physics run at 50ns, there have been some indications that the electron cloud has not yet systematically disappeared in all regions of the LHC and for all beam parameters. For instance, a pressure rise in the common beam chamber in proximity of ALICE has been observed at the locations with a wide beam chamber installed ($r = 400\text{mm}$) [5, 6]. Also, a higher heat load has been observed in some arc cells during the ramp, which has been suspected to be related to electron cloud formation from bunch shortening [7].

Machine Development with 25ns beams

Beams with 25ns spacing were injected into the LHC only during five MD sessions of the 2011 run, which are listed and briefly described here:

- (a) **29 June, 2011:** first injections of 25ns beams into the LHC. The filling scheme consisted of nine batches of 24 bunches separated by increasing gaps (2.28, 5.13 and 29.93 μs). Pressure rise around the machine as well as heat loads in the arcs were observed. All the last bunches of each batch suffered losses and emittance growth [8];
- (b) **26 August, 2011:** first injections of a 48-bunch train into the LHC with 25ns spacing. Two attempts were made to inject a 48-bunch train from the SPS, which led to beam dump triggered by large beam excursion and beam loss interlocks, respectively. During the first injection test, the transverse damper was on and it is believed that the beam suffered a coherent electron cloud




Figure 1: MD sessions labeled (a), (d) and (e): injected beams.

instability in both planes (more critical in vertical) soon after injection. During the second test, the transverse damper was switched off and the beam was affected by a coupled bunch instability [9]. This MD session had then to be interrupted because of a cryo failure caused by a thunderstorm;

- (c) **7 October, 2011:** injection tests and first ramp. In the first part of the MD, trains with 48-72-144-216-288 bunches from the SPS were injected into the LHC. Given the experience during the previous MD, the chromaticity Q' was set to around 15-20 units in both the horizontal and vertical planes in order to keep the beams stable against the electron cloud effect. In the second part, only 60 bunches per beam were injected in trains of $12 + 2 \times 24$, were accelerated to 3.5 TeV and collided during approximately 5h;
- (d) **14 October, 2011:** first long stores of 25ns beams at injection energy in the LHC. During this session up to 1020 bunches per beam were injected in batches of 72. The chromaticity was kept high in both planes ($Q'_{x,y} \approx 15$) in order to preserve the beam stability. First, a dedicated fill for pressure measurements was made, with batches injected at gradually reduced distances from 4 to 2 μs (in steps of 1 μs). Subsequently, the batch spacing was kept constant for each of the next three fills and it was set to 6.3, 3.6 and 1 μs (rounded values). Strong emittance growth and slow losses affecting the last bunches of each train were observed throughout this MD session;
- (e) **24–25 October, 2011:** record number of bunches in the LHC. Four long fills took place (average store time was approximately 4h), with 25ns beams injected into both rings in batches of 72 separated by 1 μs . In the third and fourth fills, 2100 bunches were injected for beam 1, while the number of bunches could not exceed 1020 for beam 2, due to a vacuum interlock on

one of the injection kickers (MKI). Although the situation seemed to improve over the MD, slow losses and emittance growth kept affecting both beams. Before starting the fourth fill, the horizontal chromaticity Q'_x was lowered from 15 to 3 units and the horizontal damper gain was slightly increased. Probably due to that, some horizontal instabilities could be observed from the signal of the damper pick up during the fourth fill, but the overall performance did not appear degraded from the previous fill. The MD ended with a 30' fill with only beam 1, during which batches of 72 bunches were injected into the LHC at different spacings in order to provide the stable pressure measurements needed for the modeling of the electron cloud build up in the straight sections (see next Section).

Figure 1 shows the detailed story, in terms of injected beams 1 & 2, of the sessions (a), (d) and (e). Experimental data from these MDs will be used in the next section to extrapolate the evolution of δ_{max} on the beam screen in the arcs and in proximity of the vacuum gauges. For sake of compactness, we have chosen to concatenate these three sessions and represent them as a function of a continuous time coordinate (interpretable as hours with 25ns beam), which will be systematically used throughout this paper when referring to the studies with the 25ns beams.

COMPARISON BETWEEN MACHINE DATA AND SIMULATIONS

The estimation of the δ_{max} of the chamber wall has been made separately for the straight sections and for the arcs. The former is based on the pressure data from the vacuum gauges, while the latter relies on the measured heat load on the beam screens. In reality, the SEY is not the only adjustable parameter used in our model of the secondary electron emission. Also the reflectivity of the electrons at zero energy, R_0 , and the energy at which the maximum of the SEY occurs, ϵ_{max} , are additional model parameters that

we could either fix based on previous experience or try to infer from our measurements. Laboratory measurements carried out in the past mostly assign to R_0 values ranging between 0.5 and 1 [10, 11], while a reasonable value for ϵ_{\max} for Cu surfaces seems to be around 330 eV [12].

Uncoated straight sections

The pressure measurements from the two vacuum gauges VGI.141.6L4.B and VGPB.2.5L3.B have been used for determining the pair (R_0^*, δ_{\max}^*) of the local chamber wall at a certain time. These gauges have NEG coated chambers on both sides, so that the measured pressure rise can be assumed to be dominantly linked to the electron cloud at the location of the gauge itself, whose chamber offers a baked Cu-coated surface exposed to the beam (expected initial $\delta_{\max} < 2.0$). Assuming that the desorption yield η_e is weakly dependent on the energy of the incident electrons, at least in the typical energy range covered by the LHC cloud electrons, Eq. (1) shows that there is a direct proportionality between the measured pressure rise, ΔP_{meas} , and the integrated electron flux to the wall:

$$\Delta P_{\text{meas}} \propto \int \phi_e(E) dE = \Phi_e \quad (3)$$

The procedure to estimate the model parameters R_0^* and δ_{\max}^* from pressure data and electron cloud simulations is based on the following steps [13]:

1. We inject several SPS batches into the LHC with different batch spacings and wait after each injection until the pressure reading at the gauges has leveled off to its new regime value. The pressure rise can thus be unambiguously defined at each i -th step, $\Delta P_{\text{meas}}^{(i)}$ (step 0 is used to indicate the first injection);
2. The ECLLOUD code is used to simulate the electron cloud build up for the beam configurations after each injection. A scan in the parameters (R_0, δ_{\max}) is made for each case. As a result, the total electron flux as a function of the model parameters, $\Phi_e^{(i)}(R_0, \delta_{\max})$, will be available, and also the ratios $\Phi_e^{(i)}(R_0, \delta_{\max})/\Phi_e^{(0)}(R_0, \delta_{\max})$, i.e. the fluxes normalized to that produced by the first injection;
3. We assume the pressure rise after the first injection, $\Delta P_{\text{meas}}^{(0)}$, to be our reference measurement. All subsequent measurements will be normalized to this value, $\Delta P_{\text{meas}}^{(i)}/\Delta P_{\text{meas}}^{(0)}$ in order to be compared directly with the simulated electron flux ratios and be independent of the proportionality coefficient;
4. The pair (R_0^*, δ_{\max}^*) is found as the solution of the set of equations:

$$\prod_{i=1}^n \frac{\Phi_e^{(i)}(R_0, \delta_{\max})}{\Phi_e^{(0)}(R_0, \delta_{\max})} = \frac{\Delta P_{\text{meas}}^{(i)}}{\Delta P_{\text{meas}}^{(0)}} \quad (4)$$

In principle, only two equations (or three measurements) would be sufficient to find the solution. In

practice, a higher number of equations is employed, because redundancy can resolve the ambiguity in case of multiple solutions.

The procedure above has been applied successfully four times and only with beam 1: before and after the scrubbing run, on 19 May (with 50ns beams) and on 25 October with 25ns beams (dedicated fill starting at 40h in the x-scale of Fig. 1). An attempt was also made on 14 October but, due to the rapidly changing beam and vacuum conditions, the data could not be easily fed into the procedure outlined before. The electron reflectivity at zero energy R_0^* has been found to lie in the range 0.2 – 0.3, while δ_{\max}^* has exhibited a decrease from the initial 1.9 to 1.35. The history of δ_{\max}^* is summarized in Fig. 2. In the same plot, we have also drawn the horizontal lines of the electron cloud build up thresholds at both 50ns and 25ns, as well as the vertical line representing the first injection of 25ns beams into the LHC.

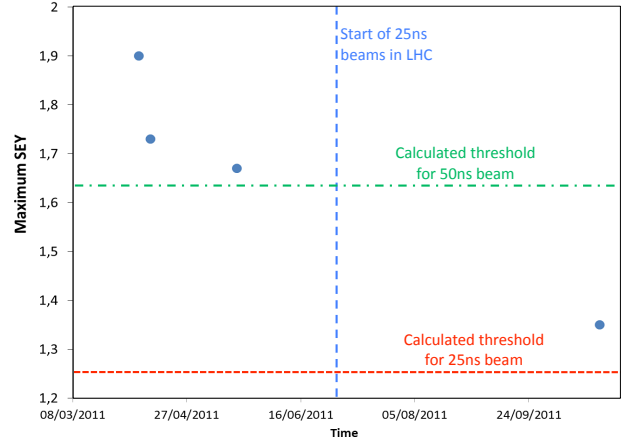


Figure 2: Evolution of δ_{\max}^* in proximity of two selected vacuum gauges from the beginning of the scrubbing run to the last 25ns run on 25 October, 2011.

It is clear that before the injection of the first 25ns beam, the SEY had only become about as low as the threshold value for 50ns beams, as is also proven by the disappearance of all the electron cloud indicators with this type of beams. Only subsequently more scrubbing has taken place thanks to the 25ns beams. The value of δ_{\max}^* has further decreased to 1.35 and is still above the build up threshold with 25ns beams. An additional substantive scrubbing step would be required to suppress the electron cloud in the uncoated locations of the straight sections.

Arcs

In the arcs, the estimation is based on the heat load data from the cryogenic system, which give the total power dissipated (in W/half-cell) on the beam screens of both beams 1 and 2. During the physics fill 1704 of 13 April, 2011, a significant heat load in the arcs (20-60 mW/m/beam) was recorded during the ramp with 228 nominal bunches per




Figure 3: Heat load measured during four fills from the MD session (d) and (e), in the same time coordinate as in Fig.1. The five vertical bars represent the measurement points used to compare heat load with electron cloud simulations.




Figure 4: Bunch-by-bunch intensity and full bunch length at the time of the fifth measurement bar in Fig. 3 (following the usual convention, blue plots on the left side refer to beam 1, while red plots on the right side are for beam 2).

beam at 50ns spacing. A comparison between these data and the power loss from the electron cloud simulated with the ECLLOUD code was carried out and is described in Ref. [14]. The outcome of the analysis was that, assuming an electron reflectivity $R_0 = 0.7$, the δ_{\max} can be estimated to be in the range 2.1–2.2, compatible with the calculated build up thresholds of 2.2 at 450 GeV and 2.1 at 3.5 TeV. More heat load observations in the arcs have been made with 25ns beams. Measurements in some reference cells from the first LHC MD with 25ns beams (MD session (a), 29 June, 2011) can be found in Ref. [8]. Figure 3 shows the heat load data, sector by sector, collected during four fills of the MD sessions (d) and (e). We can notice that the additional heat load peaked to values of nearly 50 W/half-cell (i.e. approximately an average of 0.5 W/m/beam) during the last fill with 2100 bunches for beam 1 and 1020 bunches for beam 2. A decay of the measured heat load after the last injection, and in any case after the last injection,

is also clearly visible in the examined cases and it can be explained through the weakening of the electron cloud activity due to the intensity loss (compare, for example, with the BCT signal in Fig. 1, acquired at the same time).

We know that the electron cloud in the beam chamber at a certain time, and consequently also its effects, strongly depends on the beam structure at the same time. Therefore, in order to compare the heat load data with the simulation results, we have taken five cuts in time (marked as vertical thick black lines in Fig. 3) and carried out electron cloud simulations using the correct beam structures at those times for both beam 1 and beam 2 (available from the fast BCT and Beam Quality Monitor, BQM, bunch length data). Examples of the beam snapshots in terms of bunch-by-bunch intensity and length, taken at the last measurement point of the last 25ns fill (vertical line situated furthest to the right in Fig. 3), are displayed in Fig. 4. Given the length and the irregular structures of the beams to be simulated, a new build




Figure 5: Estimated evolution of δ_{\max} on the inner surface of the beam screen in the dipole chambers.

up simulation tool had to be developed, essentially based on the same models as EPCLOUD, but with improved and optimized routines as well as capable of using the real LHC beam data as inputs. The new code has been named PyEPCLOUD, because it is almost entirely written in Python, and it has proved to be a very performant, robust and reliable tool, responding to our present needs [15].

Using the bunch-by-bunch data at the five times highlighted in Fig. 3, and also those from the MD session (a) of 29 June, PyEPCLOUD simulations were run scanning the δ_{\max} , so that the curves of the simulated heat loads $\Delta W_{b1\text{-sim}}^{(i)}(\delta_{\max})$ and $\Delta W_{b2\text{-sim}}^{(i)}(\delta_{\max})$, $i = 1$ to 6, could be produced for each i -th measurement point. The electron reflectivity at zero energy was fixed to the value of 0.7 for all simulated cases. The solution $\delta_{\max}^{*(i)}$ for the i -th heat load measurement considered is then found from the equation:

$$\Delta W_{b1\text{-sim}}^{(i)}(\delta_{\max}) + \Delta W_{b2\text{-sim}}^{(i)}(\delta_{\max}) = \Delta W_{\text{meas}}^{(i)} \quad (5)$$

By plotting the estimated $\delta_{\max}^{*(i)}$ for the six measurement points, i.e. including the measurement point during the MD session (a) plus all the five shown in Fig. 3, we obtain the evolution of δ_{\max} on the inner surface of the beam screen drawn in Fig. 5. From this graph we can see that δ_{\max} was about 2.1 when the 25ns beams were for the first time injected into the LHC. This is consistent with the aforementioned estimation based on the heat load with 50ns beams measured on the ramp on 13 April, 2011. In fact, it is reasonable to assume that no additional scrubbing could be achieved with 50ns beams between 13 April and 29 June, as the number of bunches was ramped up to 1380 per beam but the electron cloud was consistently absent in the arcs during this time. From 29 June on, with more and longer 25ns MDs, the electron cloud could be revived and, as a consequence, the LHC was further scrubbed. According to the evaluation from the last two measurement points chosen during the same store, the δ_{\max} has presently decreased down to 1.52. Looking at the last three points in the plot of Fig. 5, it is clear that, while a little scrubbing effect is still observed between the second to last and the last 25ns store, no significant decrease of δ_{\max} can be detected over the last

two points belonging to the same store. This suggests that, although the electron cloud has not yet disappeared from the arcs, the δ_{\max} has already entered a region in which the electron doses required to continue the scrubbing can be only accumulated over much longer times. The decay of the heat load after the last injection cannot be associated to a weakening of the electron cloud due to scrubbing, but only due to beam loss.

In Fig. 5 we have displayed the scrubbing curve of the arc chamber along with the values of threshold δ_{\max} for electron cloud build up with 25ns beams in the LHC at injection energy (450 GeV) and at the present top energy (3.5 TeV). Suppressing the electron cloud for this type of beams would therefore still require a decrease of δ_{\max} by approximately 0.1 at 450 GeV and 0.18 at 3.5 TeV.

Another very interesting by-product of our PyEPCLOUD simulations is the calculation of the bunch-by-bunch energy loss per turn. This is based on a simple balance on the energy of the electron cloud. The difference between the total energy of the electrons (electrostatic and kinetic) before and after the bunch passage plus the energy lost in electron-wall collisions during the bunch passage represents the net energy transferred from the bunch to the electrons — and therefore, lost by the bunch. The bunch-by-bunch energy loss per turn thus calculated can be directly compared with the one estimated from the stable phase shift measurements. This type of measurements was already applied in 2010 with 50 and 75ns beams, and then again during the scrubbing run in 2011 to qualify the efficiency of the scrubbing process [16]. However, in these cases it was always based on the global shift of the stable phase averaged over the whole beam. The method was then further refined during the 25ns MDs, when the stable phase shift could be measured in a bunch resolved fashion. This was expected to add another important piece of information to the whole picture, showing the evolution of the electron cloud along the bunch train. The data acquired at the time of the fifth measurement from Fig. 3 have been plotted together with the simulated energy loss, as resulting from the PyEPCLOUD simulation matching the measured heat load at the same time (i.e. for $\delta_{\max} = 1.52$). The resulting plot

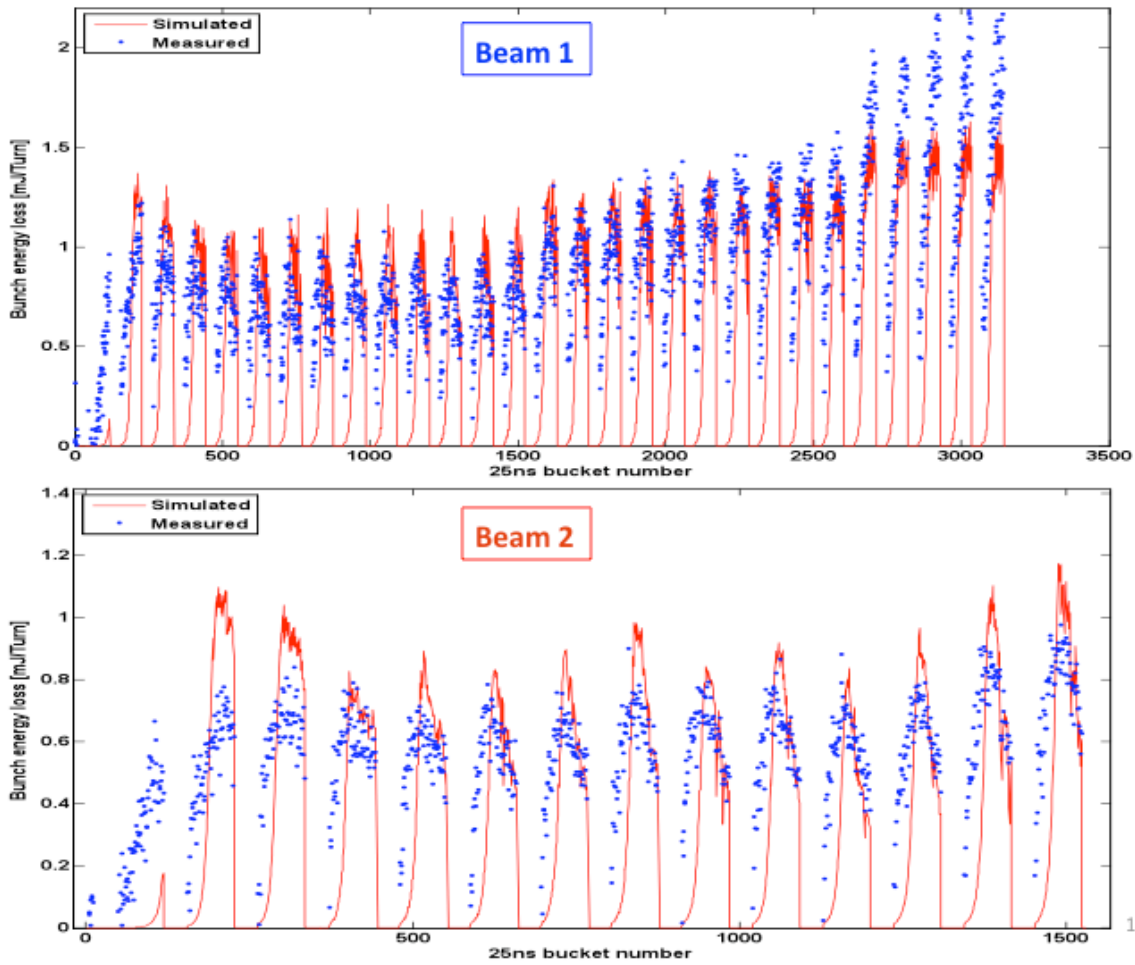


Figure 6: Bunch-by-bunch stable phase shift for beams 1 and 2: measurements and simulations.

is shown in Fig. 6. The agreement between the experimental data and the simulations is excellent, especially in the regions with saturated electron cloud. It is clear that the data tend to indicate a faster build up of the electron cloud in the machine with respect to the one predicted by simulations. This becomes especially evident if we zoom only on one batch from the full train and we look into the details of the bunch-by-bunch trend. Figure 7 shows detailed zooms on selected batches from beam 1 and 2. From this figure, we can see how the simulation is able to catch the correct values and trend of the bunch-by-bunch energy loss in the region of saturation of the electron cloud, while the evolution during the build up remains not fully captured by our model. Possible reasons for that are:

- The simulation underestimates the primary electron generation. The seed electrons come only from beam induced gas ionization, so either the dynamic pressure value or the cross section of the event might be larger than the assumed values;
- In the machine, there is a stronger memory effect between batches than that assumed in our model. This could be due either to a value of R_0 higher than that assumed in simulations (i.e. 0.7), or to the presence of uncaptured beam in the gaps, which helps the electron

survival;

- The dynamic range of the measurements cannot cover the orders of magnitude difference during the electron cloud build up;
- The measurement of the stable phase shift also includes the energy loss from other sources, e.g. impedances, which should be correctly taken into account.

In the future, we could also envisage to extract the energy loss information from the beam position data of the BQM, which provides the bunch-by-bunch position in the bucket and includes the contributions to the synchronous phase shift from both the electron cloud and beam loading. Finally, it is worth noticing that the successful comparison of the bunch-by-bunch energy loss data with the outcome of the simulation matched to the heat load data is also a non-trivial confirmation of the correct cross-calibration between the measured heat load data and the expected beam energy loss.

Beam quality

While the 50ns beam proved to be stabilized in the LHC by the electron cloud mitigation achieved with the scrub-

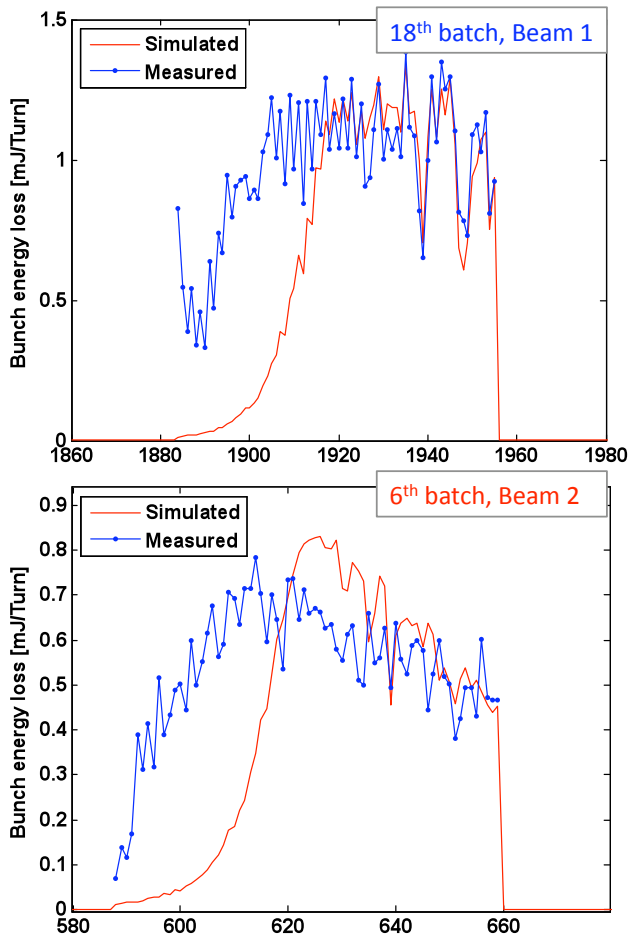


Figure 7: Close up on one selected batch of the bunch-by-bunch stable phase shifts of Fig. 6. The top plot is extracted from beam 1 and the bottom plot from beam 2. Both measurements (blue lines) and simulations (red lines) are displayed in these plots.

bing run, the 25ns beam has exhibited clear signs of transverse instability and emittance growth throughout all the dedicated MD sessions. Despite a clearly improving trend from one fill to the next one, these signs have not completely disappeared. During the first tests on 29 June, when only batches of 24 bunches were injected from the SPS, the beam could be kept inside the machine because the level of electron cloud reached along each batch was enough to cause significant emittance growth, but no coherent instability and fast beam loss [8]. When, on the following MD session, batches of 48 bunches were for the first time transferred from the SPS to the LHC, the beam was twice dumped after few hundreds of turns, due to the excitation of a transverse instability leading to unacceptable beam losses. During the successive MD sessions, this problem was circumvented by injecting the beam into the LHC with high chromaticity settings. Values of $Q'_{x,y}$ around 15 were chosen, as they had been found to be sufficiently stabilizing in HEADTAIL simulations [17]. Using these settings, the beam could be kept inside the LHC, albeit with degraded transverse emittances (see bunch-by-bunch emittance plots

from the MDs of 14, 24 and 25 October, Fig. 8). Since the BSRT needs about 2 sec to measure the emittances of each bunch, each of the snapshots in the figure does not represent an instantaneous photograph of the beam at a certain time, but results from a sweep over the bunches that can last as much as several minutes. Although the batch spacing was decreased from $2 \mu\text{s}$ during the measurement of 14 October to the $1 \mu\text{s}$ of the last MD session, the vertical emittance blow up exhibits signs of improvement. No significant further change is observed then in the vertical plane between the measurements taken in the last two fills (consistently with a slight scrubbing effect between them). The situation looks more complicated in the horizontal plane. Here a deterioration can be noticed from the 14/10 measurement to the 24/10 one. If this is related solely to the decreased batch spacing, which has enhanced the electron cloud along the full train owing to the stronger memory effect between batches, we could not explain why we observed an improvement in the vertical plane, instead. It is interesting that the situation appears improved for the 25/10 measurement, when the LHC was run with lowered hor-

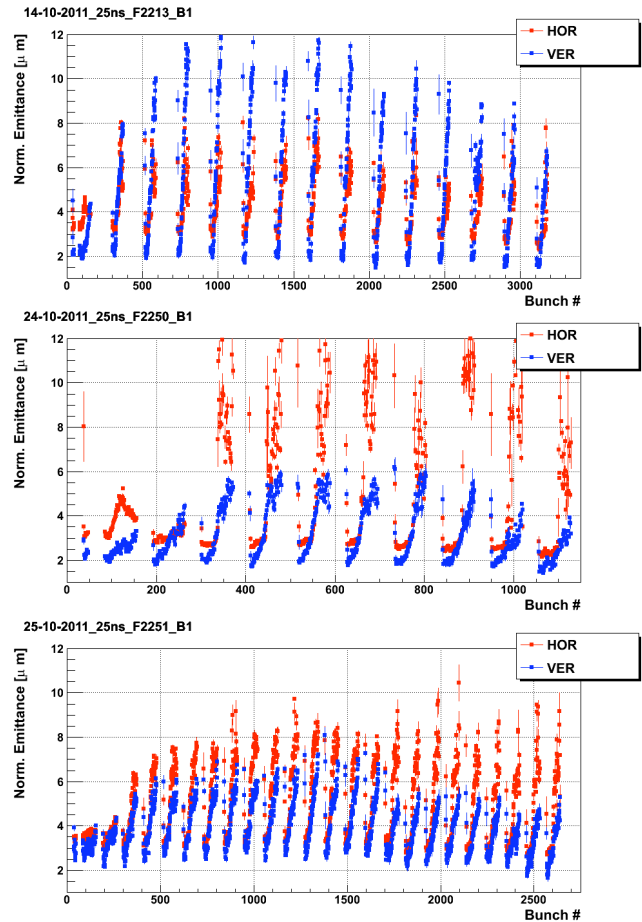


Figure 8: Snapshots of the horizontal and vertical emittance measurements for beam 1 during the last fill of 14 October and the last two fills of 24–25 October MDs.

horizontal chromaticity settings. This fact may suggest that by lowering chromaticity we have moved from a regime of strong incoherent emittance growth driven by electron cloud and high chromaticity to a new one, in which the beam suffers a fast instability, but later evolves with a better lifetime. In any case, as a general consideration, a clear weakening of the electron cloud effect from 14 to 25 October is witnessed by the improved quality of the first two–three batches. The first two seem to be hardly affected by emittance growth in both transverse planes by the time of the last 25ns fill.

The presence of electron cloud instabilities in the LHC should not be surprising with 25ns beams, as the calculated central density threshold $\rho_e = 10^{12} \text{ m}^{-3}$ in the dipoles [17] is actually always reached as long as the δ_{max} of the dipole chamber remains above the threshold for electron cloud build up. This is illustrated in Fig. 9, in which the electron cloud central density is plotted as a function of δ_{max} in the arcs and also the line of the instability threshold for nominal bunch current is drawn. A consequence of this situation is also that, unlike with 50ns beams, the margin to operate with electron cloud and stable beams is narrower with 25ns beams: to operate in a regime free of electron cloud instability, we need to either suppress the electron cloud completely or keep it in a parameter range such that it does not reach saturation for the used filling pattern.

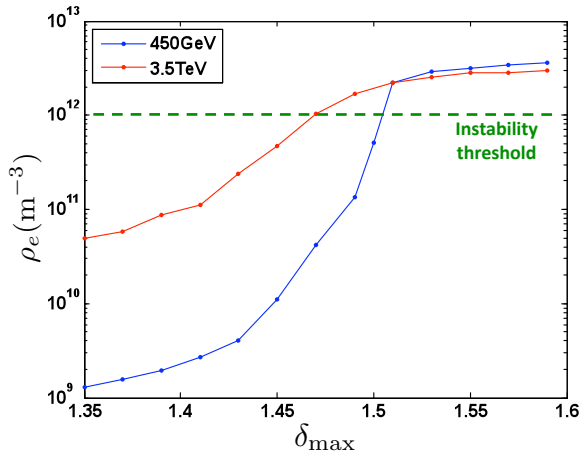


Figure 9: Electron cloud central density as a function of δ_{max} in the arcs. The line representing the instability threshold for nominal bunch current is also shown.

Lastly, other beam observables that can be used to assess the scrubbing with 25ns beams are the bunch-by-bunch beam losses and lifetimes. Losses mainly affecting the bunches at the tails of the injected batches are systematically visible in all the 25ns fills, but an improvement has been noticed over time. Figures 10 show for instance the bunch-by-bunch beam losses (top) and lifetimes (bottom) for the first three batches of the last three fills of the MD session (e) of 24–25 October, 2011. A steady improvement in the loss pattern can be observed. During the first fill the

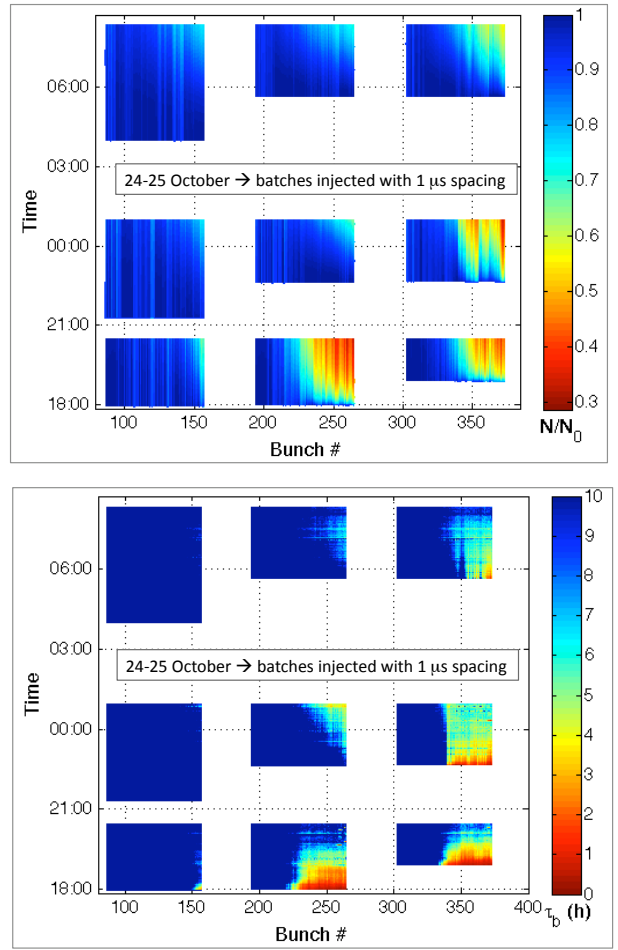


Figure 10: Bunch-by-bunch beam losses (top plot) and lifetimes (bottom plot) of the first three batches during the last three fills of the MD session (e).

first batch exhibited hardly visible signs of losses after a store of approximately 2.5h. A large fraction of the second and third batch suffered fast losses, which reduce the bunch intensity of the second half of the batch to 30% of their initial values within the first hour. During the second fill the losses on the first batch have almost entirely disappeared, while in the second batch they only become visible in our color scale after about 2h. Some bunches at the end of the third batch are still affected by fast losses, like in the previous fill. During the third fill the situation seems improved, because the losses on the third batch have become slower. Looking into the bunch-by-bunch lifetimes (essentially the derivatives of the previous plots), we also notice the existence of two different loss regimes. During the first fill, losses seem to always occur very quickly and then the bunch-by-bunch lifetimes improve as the bunches have already lost a significant fraction of their intensities. On the other hand, during the second fill we can already observe that the second batch is not affected by fast losses anymore, as it rather tends to lose at a later stage and has a deteriorating lifetime. This could be interpreted as the signature of an incoherent effect, which gently pushes the protons to-

wards the large amplitudes and the losses only occur when these particles are then scraped at the machine aperture restrictions. The third batch is still entirely in the fast loss regime. During the third fill, the bunches at the end of the second batch exhibit a deteriorating lifetime, followed by a phase of recovery. This could indicate that the mechanism underlying the incoherent effect has weakened with respect to the previous fill and protons stop being drained into a halo after the first hour of operation. The third batch is divided between an intermediate part of few bunches, which suffers from degrading lifetime (weak electron cloud), and the very last bunches that are still affected by fast losses. Hence, we can assume that by the time of the last long fill on 25 October, the δ_{\max} has reached a value such that the electron cloud fully saturates only at the end of the third batch (with a batch spacing of $1 \mu\text{s}$). This is further confirmed by the very last fill of 25 October, aimed at collecting stable pressure measurements for the estimation of δ_{\max} in the uncoated regions. During this fill, batches with variable spacing from 4 to $1 \mu\text{s}$ were injected into the LHC and, despite the pressure rise associated with each injection, no evident sign of quality degradation was observed on the beam.

SUMMARY AND CONCLUSIONS

Based on pressure data from two selected vacuum gauges and the heat load measured on the beam screen in the arcs, we have estimated the evolution of δ_{\max} over the scrubbing run with 50ns beams and over the 25ns MD sessions. The present status of the machine is summarized in Table 1, in which the threshold values for electron cloud formation with 25ns and 50ns beams are also reported.

Table 1: Estimated and threshold δ_{\max} values in the uncoated straight sections and in the arcs.

	Uncoated straight section	Arc dipoles
Estimated δ_{\max}	1.35	1.52
Threshold δ_{\max} (25ns, 450 GeV)	1.25	1.45
Threshold δ_{\max} (25ns, 3.5 TeV)	1.22	1.37
Threshold δ_{\max} (50ns, 450 GeV)	1.63	2.2
Threshold δ_{\max} (50ns, 3.5 TeV)	1.58	2.1

It was found that, before starting injection of 25ns beams into the LHC, the δ_{\max} in the different regions had reached values very close to the thresholds for electron cloud build up with 50ns beams. This allowed for a safe and stable operation with this type of beams, although some electron cloud related phenomena could still be observed in the common beam regions as well as for slight variations of relevant beam parameters (e.g. the bunch length). After the

25ns MDs, the δ_{\max} has decreased to values well below the build up thresholds for 50ns beams. The achieved values of δ_{\max} are certainly low enough as to ensure ecloud-less operation with nominal 50ns beams. PyELOUD simulations of 50ns beams with higher charges per bunch are presently ongoing to prove that stable operation should also be guaranteed in the intensity range targeted for 2012 (i.e. up to 1.8×10^{11} ppb).

Beams with 25ns spacing still produce electron cloud in the LHC, as their build up threshold values of δ_{\max} are lower than those currently achieved. Therefore, they are also affected by detrimental processes like coherent instabilities and emittance growth, which lead to fast degradation of the beam quality. Additional scrubbing is required to suppress the electron cloud with this type of beams. However, the time needed for that could be prohibitively long, because the scrubbing process slows down in time. This is because: 1) laboratory measurements show that the decrease of δ_{\max} with the bombarding electron dose has a logarithmic behaviour [12]; 2) while we are scrubbing, the electron cloud production is reduced, and so is the flux of scrubbing electrons to the walls. In other words, the electron dose required to reduce the δ_{\max} of the Cu surface of the beam screen from the present 1.52 to 1.45 is expected to be much larger than the integrated one already used to reduce the δ_{\max} from 2.1 to 1.52. Besides, this dose has to be accumulated now with a much lower electron flux. To gain speed in scrubbing, it could be envisaged to inject multi-batch trains from the SPS in future 25ns MDs in order to reach saturation of the electron cloud without requiring several 1-batch injections separated by 925ns. Another option to enhance the cleaning process could be to ramp the 25ns beam to 3.5 TeV. This would allow us to make use of the cleaning effect from synchrotron radiation, photoelectrons as well as potentially more electron cloud, given the lower build up threshold for the top energy parameters. However, beam stability could become an issue at higher energy, because running with high chromaticity is hampered by the excitation of higher order head-tail modes [2].

ACKNOWLEDGEMENTS

The authors would like to express their gratefulness to V. Baglin, P. Baudrenghien, C. Bracco, G. Bregliozzi, B. Goddard, W. Höfle, M. Jimenez, V. Kain, G. Lanza, K. Li, T. Mastoridis, H. Maury-Cuna, E. Métral, G. Papotti, S. Redaelli, B. Salvant, M. Taborelli, C. Yin-Vallgren, D. Valuch and F. Zimmermann, for their active participation in the MDs, for the support they provided with the machine operation and the experimental data, or for their relevant contributions to the simulation work. Many thanks also to W. Fischer for carefully proof-reading this paper and for his valuable comments.

REFERENCES

- [1] Proceedings of the Mini Workshop on Electron Cloud Simulations for Proton and Positron Beams, **ELOUD'02**, 15–

18 April, 2002, CERN, Geneva, Switzerland, edited by
G. Rumolo and F. Zimmermann, CERN-2002-001

- [2] N. Mounet *et al.*, elsewhere in these proceedings
- [3] G. Arduini *et al.*, [CERN-ATS-Note-2011-046 MD](#) (2011)
- [4] G. Rumolo *et al.*, [THOBA01](#) in proceedings of IPAC'11 (San Sebastian, Spain)
- [5] G. Bregliozzi *et al.*, elsewhere in these proceedings
- [6] V. Baglin, private communication
- [7] L. Taviani, "Beam Induced Heating Assessment on LHC Beam Screens" in [LHC Beam Operation Committee \(LBOC\) meeting, 27/09/2011](#)
- [8] B. Goddard *et al.*, [CERN-ATS-Note-2011-050 MD](#) (2011)
- [9] H. Bartosik and W. Höfle, [CERN-ATS-Note-2012-027 MD](#) (2012)
- [10] R. Cimino *et al.*, [Physics Rev. Lett.](#) **93**, 014801 (2004)
- [11] F. Le Pimpec *et al.*, [SLAC-TN-04-046 LCC-0146](#) (2004)
- [12] C. Yin Vallgren, Ph.D. thesis, [CERN-THESIS-2011-063](#) (2011)
- [13] O. Dominguez *et al.*, [TUPZ015](#) in proceedings of IPAC'11 (San Sebastian, Spain)
- [14] H. Maury Cuna *et al.*, [TUPZ003](#) in proceedings of IPAC'11 (San Sebastian, Spain)
- [15] G. Iadarola, "PyECLOUD", in [Electron Cloud Simulations Meeting, 28/11/2011](#)
- [16] J. Esteban-Müller, "RF Observations: Stable Phase Measurements" in [LHC Beam Operation Committee \(LBOC\) meeting, 12/04/2011](#)
- [17] K. Li and G. Rumolo, [MOPS069](#) in proceedings of IPAC'11 (San Sebastian, Spain)