Beam profile evolution induced by rf crab cavity amplitude noise in the CERN SPS

A. Fornara [*] and R. B. Appleby
The University of Manchester, Manchester, United Kingdom

H. Bartosik, X. Buffat, R. Calaga, S. Kostoglou, G. Sterbini, G. Trad, and N. Triantafyllou
CERN, Meyrin, Switzerland

(Received 6 December 2025; accepted 14 April 2026; published 8 June 2026)

In the context of the large hadron collider (LHC) luminosity upgrade, rf crab cavities (CCs) are key components for achieving the target integrated luminosity. Due to the increased bunch intensity required to achieve this goal, the collider will operate with a large crossing angle scheme to alleviate the long-range beam-beam effects. The CCs will counteract the geometric luminosity reduction factor resulting from the large crossing angle. Amplitude and phase noise injected from the CC low-level rf control are known to induce transverse bunch emittance growth. This study presents measurements performed in the CERN Super Proton Synchrotron (SPS), where two CCs were installed in 2023 for testing purposes. Only one cavity was used for this experiment. A particular focus is given to emittance growth studies driven by amplitude noise: the measured emittance growth was found to be dependent on the amplitude detuning induced by the SPS octupoles, although no dependence was predicted by the available theories and models. A new lattice model is presented that is able to reproduce both the previous measurements with phase noise and these new ones. The second part of the study is dedicated to the beam distribution evolution: the Beam Synchrotron Radiation Telescope (BSRT) measurements were able to capture the evolution of the tails of the bunch over time. Under the effect of amplitude noise, the tail contribution is reduced; in the absence of noise, it remains constant. To understand this effect, an analytical and numerical framework is used to quantify the mechanism of diffusion of the beam profile. The results of a full simulation on the SPS lattice with heavy tailed distribution evolution are then presented and discussed.

DOI: 10.1103/5wfr-3xhy

I. INTRODUCTION

After more than 10 years of operation [1], the LHC will be upgraded to the high luminosity LHC (HL-LHC) [2,3], which is expected to start beam operation in 2030. The aim is to increase the yearly integrated luminosity by a factor of 3 compared to the current LHC operational scenario, totaling more than 300 fb^{-1} per year.

To achieve this result, an increase in the bunch intensity (from 1.6×10^{11} to 2.3×10^{11} protons) and a reduction of the value of the β function at the IP, β^{*}, to 15 cm are needed. Since the LHC (and HL-LHC) operates with trains of bunches that need to share the same beam pipe to collide, long-range beam-beam effects are unavoidable [4]. This detrimental effect can be alleviated by using a large crossing angle scheme ($\frac{\theta_c}{2} = 250$ μrad), possible due to the large aperture triplet [2]. However, the introduction of this crossing angle will introduce a geometrical reduction factor $R = \frac{1}{\sqrt{1+\phi^2}}$ in the luminosity, with the Piwinski angle defined as $\phi = (\theta/2)\frac{\sigma_z}{\sigma_x}$, with σ_z and σ_x being the longitudinal and transverse beam sizes, respectively. The installation of 4 CCs [5] per beam and per IP will partially restore the loss of geometrical luminosity in the two main experiments, ATLAS [6] and CMS [7]. In addition, CCs will lengthen the longitudinal luminous region during the whole luminosity leveling process, and effectively reduce the longitudinal pile-up density of the collisions [8].

Each CC imparts a longitudinally dependent dipolar kick in the plane of the beams crossing (horizontal in ATLAS and vertical in CMS), locally introducing a correlation between the longitudinal plane and the transverse plane. The longitudinal dependence of the kick is sinusoidal in the ideal case, with a constant frequency, amplitude (given by the CC rf voltage), and phase adjusted not to kick the center of the bunch and an opposite kick to the head and tail of the bunch, producing the deflection defined as "crabbing."

[*] Also at CERN, Meyrin, Switzerland.

Fluctuations coming from the low-level rf (low-power electronics) will affect the bunch acting as a spurious source of noise. Typically, rf noise is categorized into two types: phase noise and amplitude noise. Phase fluctuations lead to a phase error of the deflecting field with the bunch core, while amplitude noise causes fluctuations in the deflecting voltage magnitude. Both types of noise induce unwanted perturbations in the CC kick, ultimately resulting in transverse emittance growth in the plane of the kick, $\frac{d\epsilon}{dt} = \dot{\epsilon} > 0$, as described in [9,10]. The maximum luminosity loss allowed due to this effect in the HL-LHC is set at 2%/h [11]. This tight constraint has driven various studies related to this topic.

The theoretical relationship between $\dot{\epsilon}$ and phase and amplitude noise in CCs has been developed in [12], and extensive studies have been carried out for the case of phase noise in [13]. In this paper, the focus is on the amplitude noise case. If $S_{\Delta A}$ [Hz^{-1}] is the relative power spectral density (PSD), normalized to the carrier PSD, of the amplitude noise across all synchro-betatron lines, we can obtain the following estimates for the emittance growth [12]:

$$\dot{\epsilon} = \beta_{\mathrm{CC}} \left(\frac{eV_{\mathrm{CC}} f_{\mathrm{rev}}}{2E_b} \right)^2 C_{\Delta A} \sum_{k=-\infty}^{+\infty} S_{\Delta A}[(k \pm \bar{\nu}_b \pm \bar{\nu}_s) f_{\mathrm{rev}}], \tag{1}$$

where β_{CC} is the β function at the location of the CC, e is the charge of the proton, V_{CC} is the voltage of the CC, f_{rev} is the revolution frequency of the bunch, E_b is the energy of the bunch, and $\bar{\nu}_b$ and $\bar{\nu}_s$ are the means of the betatron and synchrotron tune distributions. The factor $C_{\Delta A}$ is introduced to take into account the bunch length and is defined as [12]:

$$C_{\Delta A}(\sigma_\phi) = e^{-\sigma_\phi^2} \sum_{l=0}^{+\infty} I_{2l+1}(\sigma_\phi^2) \tag{2}$$

with σ_ϕ the rms bunch length (in radians) with respect to the CC frequency, f_{CC}, and $I_n(x)$ the modified Bessel function of the first kind. The $C_{\Delta A}(\sigma_\phi)$ is a monotonic function equal to 0 for a bunch length $\sigma_\phi = 0$, which increases with σ_ϕ to a maximum of ~0.25. In the experimental settings, we discuss in the following, $C_{\Delta A} \sim 0.23$, Fig. 1.

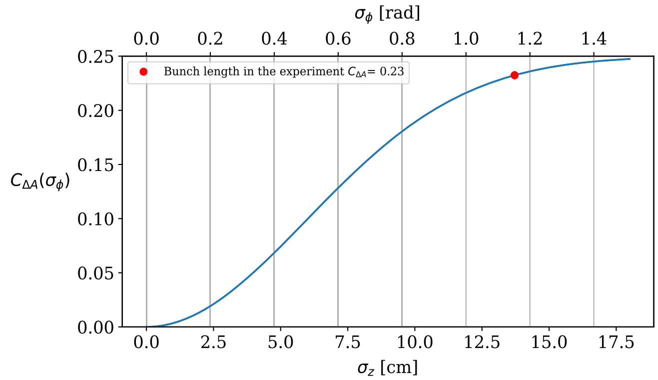

FIG. 1. $C_{\Delta A}$ function. The red dot represents the parameter value for the experimental conditions under study.

The series in Eq. (1) does not diverge for two concurrent reasons: (i) A CC is a resonator, therefore the PSD decreases rapidly away from its fundamental frequency $f_{\mathrm{CC}} \sim 400$ MHz; in particular, in our case, the injected noise falls by several orders of magnitude after 10 kHz; and (ii) The bunch samples the noise only at the synchro-betatron lines.

Following these two points, we can calculate that the whole series is reduced to four finite terms: $S_{\Delta A}(0 + \bar{\nu}_b + \bar{\nu}_s)$, $S_{\Delta A}(0 + \bar{\nu}_b - \bar{\nu}_s)$, $S_{\Delta A}(0 - \bar{\nu}_b + \bar{\nu}_s)$, $S_{\Delta A}(0 - \bar{\nu}_b - \bar{\nu}_s)$. These four points correspond to the first synchro-betatron sidebands in the machine. We note that the sum over all integers k with both sign combinations already accounts for positive and negative frequency sidebands symmetrically; for each k, the four terms listed above correspond to physically distinct synchro-betatron lines so that no double counting occurs. The value of each of these contributions can be calculated as follows:

$$(\pm \bar{\nu}_b \pm \bar{\nu}_s) \times f_{\mathrm{rev}} = (\pm 0.18 \pm 0.0051) \times 43.38 \text{ kHz}. \tag{3}$$

Therefore, the synchro-betatron sidebands are, explicitly: $(0.18 + 0.0051) \times 43.38 = 8.03$ kHz and $(0.18 - 0.0051) \times 43.38 = 7.59$ kHz, giving sidebands at (± 8.03 kHz and ± 7.59 kHz). Now, as visible in Fig. 2, the noise is flat in the region comprising the sidebands, therefore, in our specific case Eq. (1) can be written as

$$\dot{\epsilon} = 4\beta_{\mathrm{CC}} \left(\frac{eV_{\mathrm{CC}} f_{\mathrm{rev}}}{2E_b} \right)^2 C_{\Delta A}(\sigma_\phi) \bar{S}_{\Delta A}, \tag{4}$$

where $\bar{S}_{\Delta A}$ is any of the four finite terms.

A. Effect of the beam coupling impedance on $\dot{\epsilon}$

The analytical model Eq. (1) for amplitude noise in [12] describes an induced increase in emittance through decoherence mechanisms. If a dipolar noise is applied to an ensemble of particles without tune spread, then all the particles will be displaced by the same amount, creating a coherent oscillation of the bunch. Instead, if a betatron tune spread is present due to nonlinearities, phase mixing between particles within the bunch leads to decoherence of the coherent betatron oscillations induced by noise. This subsequently leads to emittance growth. However, the analytical model does not include the effect of the beam coupling impedance in the analytical estimate of $\dot{\epsilon}$ given by Eq. (4) [8].

The experimental campaign on phase noise [13] showed a dependence of the measured $\dot{\epsilon}$ on octupole strength through a mechanism based on the overlap between the

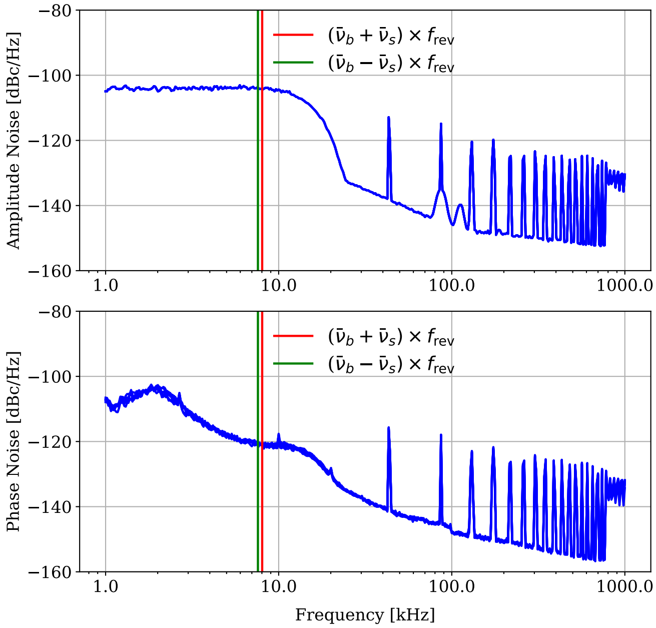

FIG. 2. Amplitude and phase noise measurements from the signal noise analyzer, taken with the beam circulating. Phase noise is around 20 dB lower than the amplitude noise near the main lines (red, green). As a reference, following [12], a phase noise at $-120\,\mathrm{dBc/Hz}$ corresponds to an emittance growth of around $0.5\,\mu\mathrm{m/h}$, while an amplitude noise at $-104\,\mathrm{dBc/Hz}$ corresponds to around $20\,\mu\mathrm{m/h}$. The peaks visible at higher frequencies are f_{rev} harmonics arising from the direct beam-induced voltage in the cavity.

coherent betatron tune and the incoherent tune spread distribution.

In Ref. [14], it was shown that the $\dot{\epsilon}$ induced by the CC phase noise can be suppressed by the beam coupling impedance, when the mode 0 (dipole oscillation) is naturally damped by the impedance and its tune shift is pushed outside of the incoherent spectrum, a condition that is valid for the SPS CC experiments. In short, the transverse impedance generates a real tune shift that pushes the coherent mode outside the incoherent spectrum. In this regime, the response to the noise kick is dominated by the coherent mode, which is damped by the imaginary part of the impedance-induced tune shift. Since only the incoherent spectrum contributes to emittance growth [14], and its excitation is suppressed when the coherent mode dominates, $\dot{\epsilon}$ is reduced below the prediction of Eq. (4). When sufficient octupole-induced tune spread is added, the incoherent spectrum broadens to overlap with the coherent tune, the coherent mode is reabsorbed into the incoherent distribution, and $\dot{\epsilon}$ recovers the value predicted by Eq. (4). This condition was reproduced in tracking simulations [13] by incorporating the SPS transverse impedance model in PyHEADTAIL [15,16].

With amplitude noise, a mode 1 (head-tail) motion is induced. As opposed to the mode 0, this motion is naturally enhanced by the impedance in the experimental conditions. A suppression of the emittance growth was therefore not expected and, indeed, in [13] the simulations do not predict a dependence on the tune spread.

In the first part of this work, the aim is to complement these previous studies, which mainly focused on phase noise, through the experimental study of CC amplitude noise in the SPS. The aim is to test the model of Eq. (1) and to understand whether it is capable of describing the experimental results in the case of amplitude noise. A more refined lattice simulation for the SPS with a CC installed is then presented and compared to both the old and new datasets.

II. SPS EXPERIMENT

A. General setup

The experimental setup is the same one proposed for phase noise studies [13]: a bunch with 3×10^{10} protons stored at a constant energy of 270 GeV in the SPS machine. The optics used has $(Q_x, Q_y) = (26.13, 26.18)$ [17]. A single vertical CC is active during the measurement campaign, with the cryomodule containing the CC installed on a mobile table that allows it to be moved into the beam line as needed during the MD and to be removed for normal operations. The CC beam pipe is isolated from the SPS vacuum by a gate valve, so moving the cryomodule in and out of the beam line does not require breaking the machine vacuum.

During the experiment, two main nonlinearities are present: the sextupoles to provide a positive chromaticity value of $Q' \approx 1$ in order to stabilize the beam, and the Landau octupoles for amplitude detuning (referred to as Landau octupole defocusing, LOD from now on) [18]. The LOD circuit is present at positions with large vertical beta-function to provide bunch stabilization on the vertical plane. The current in the LOD circuit was varied to investigate the effect of different detuning values (corresponding to the octupolar strengths K_{LOD} in the range $[-30\,\mathrm{m}^{-4}, 30\,\mathrm{m}^{-4}]$ [19], where K_{LOD} denotes the integrated octupole strength as defined in MAD-X [19]). A summary of the fundamental parameters is shown in Table I.

PSD was measured by a signal noise analyzer, which provides the single sideband measurement with respect to the carrier (the CC resonant frequency) in dBc/Hz to be converted to Hz^{-1} following the definitions in [20]. In the experimental condition, the PSD was -104 dBc/Hz, that is $4\times 10^{-11}\ \mathrm{Hz}^{-1}$. During the measurement, the phase noise level was maintained at about 20 dB lower than the amplitude noise level in the frequency range of interest, to ensure an amplitude noise dominated regime. Following the scaling of Ref. [12], a phase noise at -120 dBc/Hz corresponds to $\dot{\epsilon} \approx 0.5\,\mu\mathrm{m/h}$, while the amplitude noise at -104 dBc/Hz corresponds to $\approx 20\,\mu\mathrm{m/h}$. The phase noise contribution is therefore about $\sim 2.5\%$ of the total measured growth rate, and the two contributions can be treated as

TABLE I. Lattice specifications.

Parameter	Value
(Q_x, Q_y)	(26.13, 26.18)
(Q')	1
Bunch intensity	3×10^{10}
Bunch length $4\sigma_t$	1.8 ns
Initial ϵ_x, ϵ_y (normalized)	$\sim$3 μm
K_{LOD}	$[-30\ \mathrm{m}^{-4}, 30\ \mathrm{m}^{-4}]$
α_{yy} at $\lvert K_{\mathrm{LOD}}\rvert = 10\ \mathrm{m}^{-4}$	$\sim 10^4\ \mathrm{m}^{-1}$
V_{CC}	1 MV
f_{CC}	400 MHz
$(\beta_{\mathrm{CC},x}, \beta_{\mathrm{CC},y})$	(30 m, 78 m)

effectively separated. The transverse feedback damper was not active during the measurements. Given the low bunch intensity of 3×10^{10} protons per bunch, the beam was sufficiently stable without it. The noise data, collected from the direct rf signal coming from the field antenna, which measures the voltage in the cavity, can be seen in Fig. 2.

B. Measurements

In this section, the experiment will be described. A given amplitude noise level is injected in the CC low-level electronics for about 10 min, enough to measure the emittance growth, while the low-intensity bunch is circulating in the SPS with constant energy. During this time, the emittance is measured every $\sim$30 s. The bunch is then extracted, and a new measurement begins. During previous studies, wire scanners [21] were used to obtain transverse distribution information, but at the time of this study they were not available [22]. The Beam Synchrotron Radiation

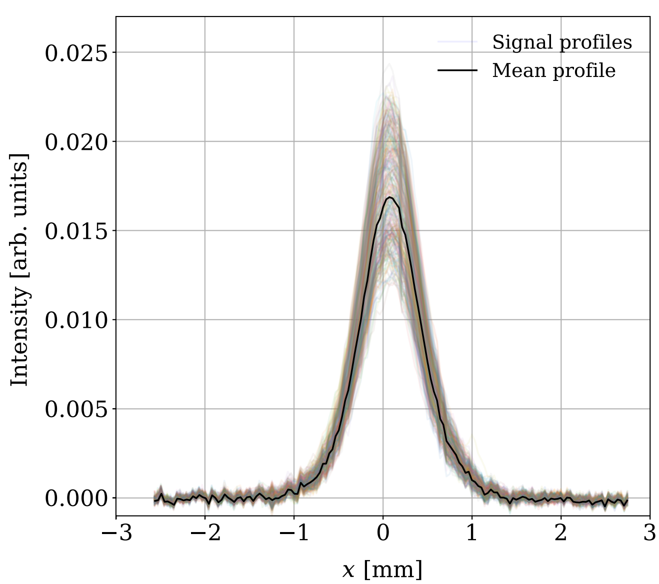

FIG. 3. BSRT profiles measured during a single acquisition. A Gaussian fit is performed to each profile and the average variance is extracted.

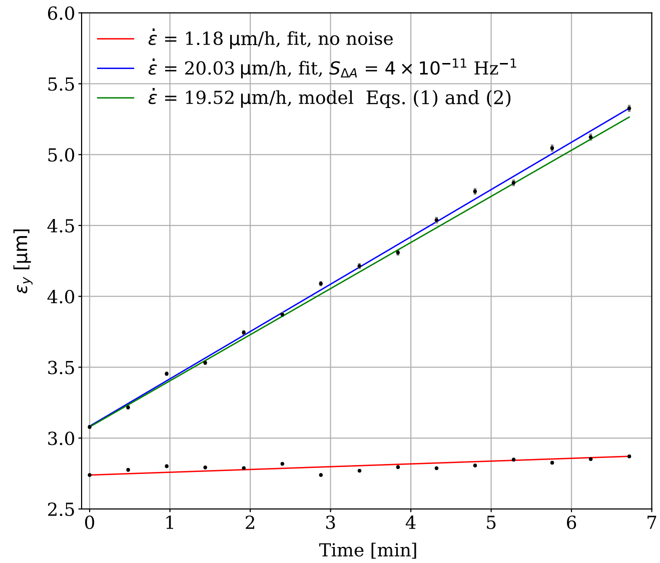

FIG. 4. Data and corresponding linear fit to the emittance evolution during two different fills, the first one with no noise injected (red) and the second one with a $S_{\Delta A} = 4 \times 10^{-11}\ \mathrm{Hz}^{-1}$ (blue). The second case is then compared with Eq. (1) in (green).

Telescope (BSRT) [23] was used as a diagnostic device to characterize the vertical bunch profile. A dedicated BSRT configuration was used to minimize the uncertainty in the measured ϵ_y: a set of several hundreds of turn-by-turn transverse measurements of the bunch is provided, which can then be fitted with a Gaussian distribution and averaged to obtain ϵ_y at a particular time. These consecutive measurements are taken within several microseconds, a relatively small window to assume that the emittance remains constant during the measurement time, as shown in Fig. 3. Several of these measurements are then used to perform a linear fit to $\epsilon(t)$ to extract the average $\dot{\epsilon}$.

The first set of measurements is dedicated to an estimation of the residual emittance growth if there is no injected noise in the machine. The other sets of measurements are instead related to fills with amplitude noise injected in the CC. In Fig. 4, the average $\dot{\epsilon}$ in the case of $S_{\Delta A} = 4 \times 10^{-11}\ \mathrm{Hz}^{-1}$ is shown together with the first noise-free measurement.

Two things stand out in Fig. 4: first, the high precision of the BSRT in measuring the emittance growth. Second, there is very good agreement between the model and the measured values. So, thanks to its precision, a large set of measurements, and negligible vertical dispersion at its location, the BSRT was able to measure the evolution of the beam transverse distribution quite accurately.

III. EMITTANCE GROWTH SIMULATIONS

In this section, the emittance growth and its dependence on the amplitude detuning present in the machine during the measurement are analyzed. The measurements carried

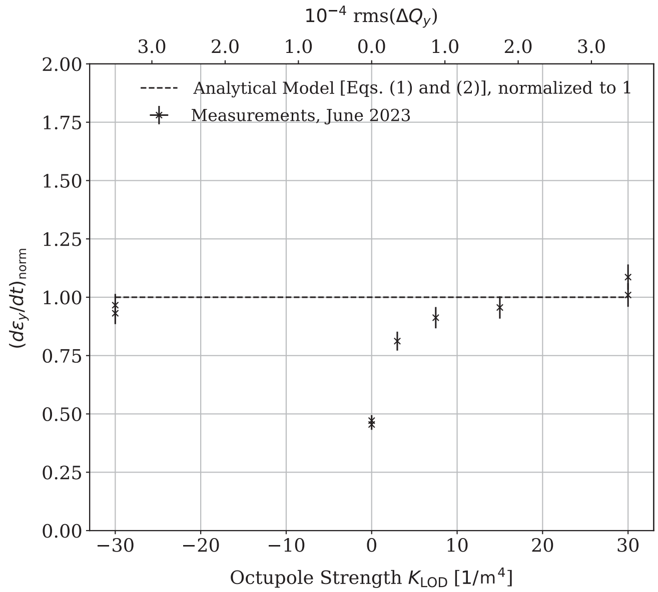

FIG. 5. Measured emittance growth rate, normalized to the analytical prediction of Eqs. (1)–(4), as a function of the octupolar strength in the LOD circuit. Data points (black) reveal a dependence on the octupolar detuning that is not predicted by either the zero coupling model or the simulations in [13].

out in the SPS in 2023 and discussed in the previous section are summarized in Fig. 5. A clear dependence of $\dot{\epsilon}$ induced by amplitude noise in the CC is observed as a function of the octupole strength (and thus of the incoherent tune spread). This dependence was not expected either in Eq. (4) or from the simulation model previously developed in [13].

In order to investigate this observed phenomenon, a dedicated simulation model based on a thin lens SPS lattice has been developed and implemented using the tracking tool Xsuite [24–26].

A. Simulation setup

Collective effects coming from the beam coupling impedance were included in the simulations by leveraging the PyHEADTAIL interface available in Xsuite, which allows the definition of wakefields and the calculation of their turn-by-turn effect on the bunch. By varying the octupole strength K_{LOD} in the range $[-30\ \mathrm{m}^{-4}, 30\ \mathrm{m}^{-4}]$ [17] in the simulations, the scan performed during the measurement is reproduced in the model. Note that in the experiments and new simulations, the octupoles are distributed along the line and are not taken into account as a single detuning term according to the model in [13]. A CC element is installed in the lattice, where $(\beta_{\mathrm{CC},x}, \beta_{\mathrm{CC},y}) \sim (30\ \mathrm{m}, 78\ \mathrm{m})$. The CC has a voltage of $V_{\mathrm{CC}} = 1$ MV, and a frequency of $f_{\mathrm{CC}} \sim 400$ MHz.

The CC kick with phase and amplitude noise at turn N is modeled as follows:

$$\Delta p_y(N) = \frac{eV_{\mathrm{CC}}}{E_b}[1 + \Delta A(N)] \times \sin\left[\frac{2\pi f_{\mathrm{CC}}}{c\beta_0} z + \Delta\phi(N)\right]. \tag{5}$$

The parameters ΔA and $\Delta\phi$ are random numbers drawn from a Gaussian distribution with $\mu = 0$ and variance:

$$\sigma^2_{A,\phi} = S_{\Delta A,\Delta\phi} \times f_{\mathrm{rev}}. \tag{6}$$

This assumption is justified by measurements from the spectrum analyzer that show a flat noise distribution up to well beyond the first betatron line, see Fig. 2. In the following, either phase or amplitude noise is used in a given calculation, so the other is set to zero.

Three main sources of transverse impedance [27] are taken into account: resistive wall impedance, SPS kickers, and steps of the vacuum chamber transition. Both dipolar and quadrupolar components [28] of the impedance are taken into account. Studies including the impedance of the main fundamental mode of the CC were also performed, showing that there was no significant effect for the purpose of this analysis. In the model, each source is included as an element in the lattice line that acts on the bunch via a kick that decays over several turns and is weighted by the β-function at the location of each source.

In this section, related to emittance growth, the bunch has a Gaussian distribution in all transverse coordinates and is matched longitudinally to the nonlinear rf bucket. The emittance is calculated as

$$\epsilon_{\mathrm{rms},y} = \sqrt{\langle y^2\rangle\langle p_y^2\rangle - \langle y \cdot p_y\rangle^2}, \tag{7}$$

where y and p_y are the coordinates of each particle with its own dispersion taken into account.

The $\dot{\epsilon}$ obtained in a single simulation is subject to numerical fluctuations due to the limited amount of particles that can be tracked in a reasonable amount of time. To solve this problem, multiple parallel simulations are performed on GPU on the same lattice with different sets of particles, and the average between different simulations is taken as a measure of $\langle\dot{\epsilon}\rangle$. In this study, $\mathcal{O}(100k)$ particles per simulation are tracked for $\mathcal{O}(100k)$ turns, and each simulation is repeated 10 times to obtain the average and standard deviation. All these simulations are performed on the CERN High-Throughput Computing Batch System [29] based on HTCondor [30]. Each simulation takes $\mathcal{O}(1\ \mathrm{h})$ to complete.

The first step of the study was to benchmark the new simulation against previous measurements [13] in relation to phase noise, showing that the new framework is capable of predicting the outcome of previous experiments (the Appendix). After this benchmark, it was possible to apply the new method to study the data from 2023 related to amplitude noise.

B. Results

The results of the Xsuite simulations are reported in Fig. 6 together with the experimental measurements. The results obtained with the new simulation model agree very well with the measurement data and show the observed dependence of the emittance growth induced by CC amplitude noise on the tune spread. However, the dependence of this growth on the octupole strength is not the same as that of phase-noise-induced emittance growth. In particular, in the case of amplitude noise, the prediction of Eq. (4) is restored by setting $|K_{\mathrm{LOD}}| \gtrsim 10\ \mathrm{m}^{-4}$, whereas in the case of phase noise a greater $|K_{\mathrm{LOD}}|$ is needed.

The cause for the partial suppression of the emittance growth with low octupole field in this experiment is believed to still be linked to the damping of mode 0 by wakefields, as in the case of dipole noise. However, one must consider that the beam motion is no longer purely rigid when chromaticity is included [31]. As a result, the energy from the kick is distributed to different head-tail modes depending on the chromaticity [32]. With a small positive chromaticity, as in the experiment described here, the nonrigid mode 1 remains the most impacted by the crab cavity amplitude noise. However, part of the energy is imparted to the nonrigid mode 0 and can be suppressed by its damping, exactly as in the experiments described in [13]. In addition, one may expect that the nonlinearity of the rf force further distorts the nonrigid head-tail mode and thus impacts the suppression mechanism. This second aspect is much less covered in literature, and an analytical description is beyond the scope of this paper. We therefore propose a numerical investigation of these two possible sources of discrepancy between the previous simulations and the new measured data, namely the chromaticity difference and the nonlinear longitudinal mapping, as presented in Fig. 7.

In the first campaign of measurements, $Q' \approx 0.5$ was used, while in our case $Q' \approx 1$. The lower Q' could reduce the detuning in the vicinity of the region with a low octupolar detuning, leading to a reduction in $\dot{\epsilon}$. In the previous simulations [33] in PyHEADTAIL, a linear longitudinal map is used, whereas in the experimental case the accelerating rf curvature is visible to a non-negligible portion of the bunch. To test this hypothesis, the voltage of the main rf can be doubled and its frequency halved in the Xsuite simulation to linearize the map at constant $\tilde{\nu}_s$, Eq. (1), keeping $Q' \approx 1$.

The results of this study are shown in Fig. 7.

The study shows that a linearization of the longitudinal motion cancels out the dependence for $|K_{\mathrm{LOD}}| \sim 0\ \mathrm{m}^{-4}$.

To summarize this section, it was demonstrated that the hypothesis of a suppression of the CC rf noise-induced emittance growth that was valid for phase noise is also experimentally observed for amplitude noise with the new data (Fig. 5). This dependence was unexpected and the previous models were not able to predict this feature. The model developed in Xsuite is perfectly able to reproduce the experimental data related to both phase and amplitude

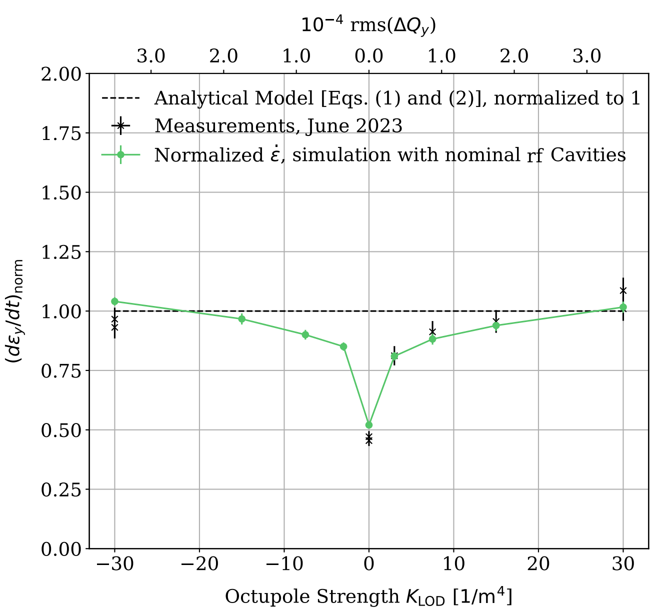

FIG. 6. Measured emittance growth rate, normalized to the analytical prediction of Eqs. (1)–(4), as a function of the octupolar strength in the LOD circuit. The new Xsuite simulation is in very good agreement with the measured data.

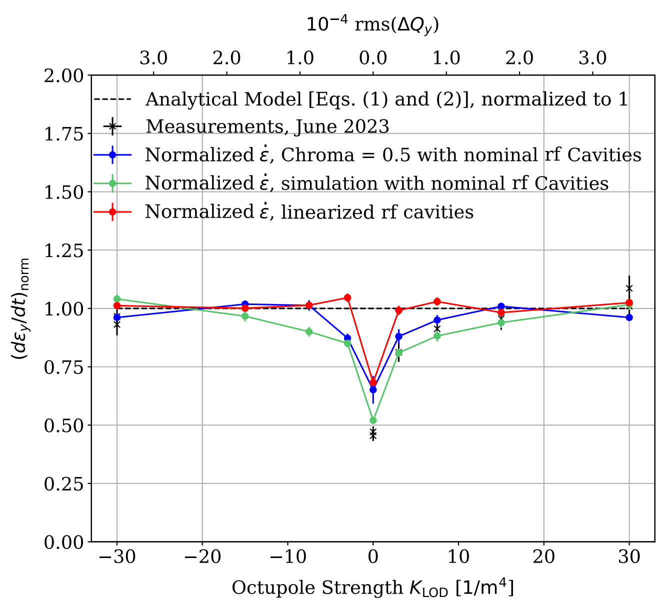

FIG. 7. Measured emittance growth rate, normalized to the analytical prediction of Eqs. (1)–(4), as a function of the octupolar strength in the LOD circuit. The simulation with decreased chromaticity (blue) shows the same characteristics as the one with the nominal value (green). The linearization of the rf cavities (red), instead, leads to the complete absence of a dependence on the octupolar strength away from 0, as shown in [13].

noise. The main source of discrepancy has also been identified as the linearization of the longitudinal motion in the old model, whereas in the model presented here the curvature of the main rf is also included.

IV. TAIL EVOLUTION

The second important part of this study is related to the evolution of the profile in the presence of CC noise during measurement. It is well known that under the effect of an external white noise, the phase space population must grow to larger amplitudes, leading to emittance growth [10]. This new set of measurements allows us to not only experimentally evaluate how the core of the bunch evolves but also how the tail population changes under the effect of amplitude noise injected in the CC. In this section, a characterization of the transverse distribution is proposed and subsequently applied to the available data. The evolution of the tails is then obtained from the dataset and a dedicated set of simulations with the SPS lattice model is developed and presented.

A. q-Gaussian and the diffusion equation

To characterize the transverse distribution, the q-Gaussian distribution [34] is introduced. This distribution is adopted in the context of accelerator physics [35,36] to parametrize distributions with heavier tails compared to a Gaussian distribution. The q-Gaussian is defined as

$$f(x) = \frac{\sqrt{\mathrm{B}}}{C_q} e_q(-\mathrm{B}x^2), \tag{8}$$

with the q-exponential defined as

$$e_q(x) = [1 + (1-q)x]_+^{\frac{1}{1-q}}, \tag{9}$$

where $\mathrm{B} > 0$ and the notation $[\cdot]_+ \equiv \max(\cdot, 0)$ ensures that the q-exponential is defined only where its argument is positive. The normalization factor is defined as

$$C_q = \begin{cases} \frac{2\sqrt{\pi}\,\Gamma\left(\frac{1}{1-q}\right)}{(3-q)\sqrt{1-q}\,\Gamma\left(\frac{3-q}{2(1-q)}\right)} & \text{for } -\infty < q < 1, \\ \sqrt{\pi} & \text{for } q = 1, \\ \frac{\sqrt{\pi}\,\Gamma\left(\frac{3-q}{2(q-1)}\right)}{\sqrt{q-1}\,\Gamma\left(\frac{1}{q-1}\right)} & \text{for } 1 < q < 3. \end{cases} \tag{10}$$

with Γ being the Gamma function. The Gaussian distribution can be retrieved in the limit $q \to 1$. In the following paragraphs, the time evolution of a q-Gaussian from $q \neq 1$ to $q = 1$ will be referred to as "Gaussian shaping." The main reason to extensively apply this distribution to the beam profiles comes from this useful property of having a single parameter that, with continuity, returns either heavy tailed, regular Gaussian, or low tailed distributions. Another important property of this distribution is shown in [37]: for a given q-Gaussian distribution in x one can apply an inverse Abel transform to it to obtain the distribution in $J = \frac{1}{2}(x^2 + p_x^2)$. The distribution in J follows a q-exponential distribution. The problem of tail population evolution under the effect of noise has been studied in [38,39]. Under the general principle of averaging over a fast varying angle variable [40] the diffusion equation for J can be written as

$$\frac{\partial \rho(J,t)}{\partial t} = \frac{1}{2}\frac{\partial}{\partial J}\left[D(J)\frac{\partial \rho(J,t)}{\partial J}\right], \tag{11}$$

where $D(J)$ is the diffusion coefficient of the problem. Note that the factor of $\frac{1}{2}$ follows the convention of Refs. [38,39]; replacing $D(J)$ by $2D(J)$ would recover the standard form of the Fokker-Planck diffusion equation, but both conventions are equivalent. In the case of dipolar white noise, the diffusion coefficient is the following [38]:

$$D(J) = \beta_D \sigma_\theta^2 f_{\mathrm{rev}} \times J = \alpha \times J, \tag{12}$$

where β_D is the β function at the dipole location and σ_θ^2 is the variance of the noise kick. Therefore, the diffusion coefficient is a linear function of the action. To demonstrate that under this condition a "Gaussian shaping" is expected, an initial q-exponential distribution is assumed for $\rho(J,t)$ and Eq. (11) is iteratively solved for an arbitrary value of α. We note that this result can also be confirmed analytically: inserting $\rho(J,t) = a(t)\exp(-a(t)J)$ into Eq. (11) with $D(J) = \alpha J$ yields the correct equation for $\dot{a}(t)$, confirming that Gaussian shaping (i.e., convergence toward a pure exponential in J) is the exact long-time behavior under this diffusion process. The numerical time evolution of $\rho(J,t)$ under Eq. (11) is shown in Fig. 8 by fitting a q-exponential to the J distribution at each time step t.

It is thus predicted by Eq. (11) that the effect of dipolar noise is to reduce the contribution of the tail population while causing emittance growth. In particular, once the value $q = 1$ is reached, the emittance continues to grow, but the tails do not change.

Since the CC gives a dipolar kick to the bunch in the transverse coordinate, this same model is capable of predicting that in the measurements and simulations a reduction of q should be present. The following part of this study aims to study in detail the evolution of the bunch tails during our measurement together with dedicated tracking simulations in the SPS lattice described before.

B. Beam profile evolution

The BSRT measurements are now used to study the evolution of the beam profile in the experiment. In this case, a q-Gaussian fit is applied to each average profile to obtain information related to the tail population evolution. After choosing one particular dataset among the several acquired

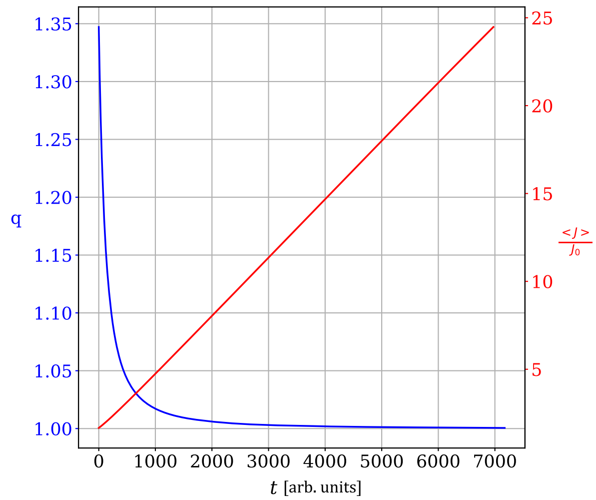

FIG. 8. Evolution of a q-exponential fit to the numerical solution of Eq. (11). The starting distribution is heavy tailed with $q = 1.35$. The subsequent $\rho(J, t)$ are subject to emittance growth and are brought a simple exponential shape under the effect of $D(J) = \alpha \times J$.

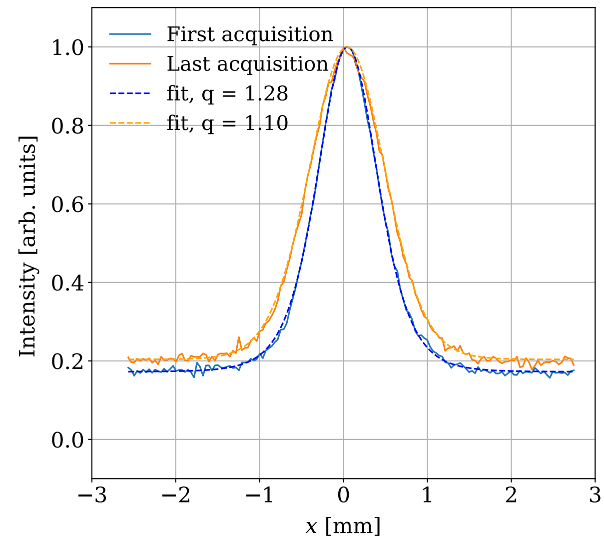

FIG. 9. Average measured profile under the effect of noise, initial and final acquisitions of the same set. The emittance of the final acquisition is increased as expected due to the effect of noise, while the q-parameter decreases.

while the bunch was subject to noise, the first and last acquisitions are plotted in Fig. 9, together with the corresponding q-Gaussian fit.

The bunch is q-Gaussian when injected, with $q \sim 1.3$. After injection, the dynamics is clearly driven by the noise in the CC, showing that the beam is subject to "Gaussian shaping." This leads to extending the study to the various measurements taken under the effect of noise, together with the only dataset without noise taken as reference in Fig. 10.

In the plot, the only datasets shown are those acquired with a $|K_{\mathrm{LOD}}| > 0$, where the contribution coming from the impedance-driven coherent motion described in the previous section is reduced. The parameter q decreases as expected from Eq. (11). The evolution of the q-parameter is a clear indication of how the transverse distribution is changing due to the noise injected: when noise is not present, there is no sensible effect, for the timespan under consideration, that is capable of changing the tail contribution to the full distribution. Once noise is injected, this behavior clearly changes, leading to "Gaussian shaping." Even though this does not mean that the only process that can lead to a diffusion of this kind is noise, in the experimental conditions under study, noise is the only driving force that could explain such a behavior in a time interval of less than 10 min.

In the next section, the transverse tail evolution of the bunch under the effect of noise is studied. To generate the q-Gaussian, a Monte Carlo approach is used as in [34].

C. SPS simulation

The tail population study in the SPS is carried out with the same map used in the emittance growth study, with the parameters extracted from Table I, a $K_{\mathrm{LOD}} = 30\ \mathrm{m}^{-4}$, and including the impedance model. Although the level of noise in the experiment was very high compared to the requirements needed for HL-LHC, both emittance growth and tail evolution required minutes of machine time to be measured. This is reflected in a much more complex simulation

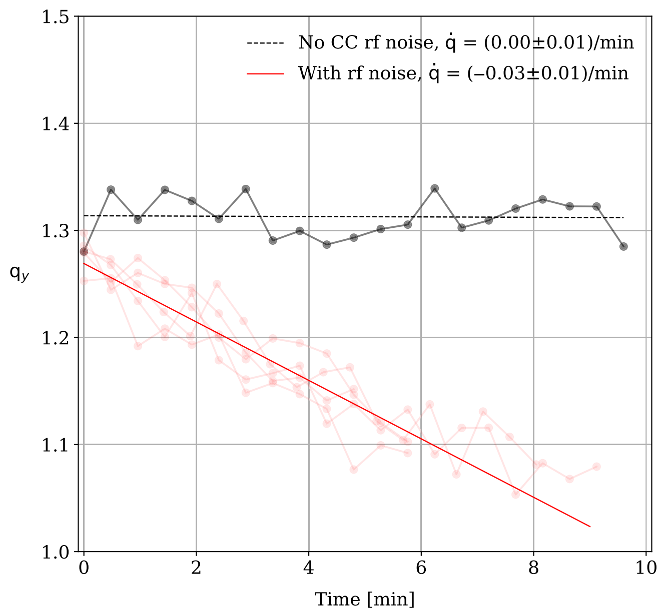

FIG. 10. q-parameter evolution as a function of time for the dataset. When noise is not injected in the CC (black) the tail population remains almost unchanged. If noise is injected (red) the bunch experiences a "Gaussian shaping" process. In the plot the average linear fit between all the datasets with noise and $|K_{\mathrm{LOD}}| > 0$ is shown, together with the statistical uncertainty on the angular coefficient of the fit.

setup for the second case, compared to the first, which required $\mathcal{O}$(1M) turns to obtain the results shown here.

To then have a statistical error small enough to compare the computational result with the data, 10 simulations are used to extract the average and standard deviation of the evolution of the q-parameter. Each simulation runs in parallel on a different GPU and requires ∼24 h of computation time. The evolution of the q-parameter in the horizontal plane (unaffected by noise) is also shown to further demonstrate the hypothesis of a noise-driven q-parameter evolution. The results of the simulation are reported in Fig. 11, to be compared with Fig. 10.

After taking the average between the 10 simulations, a linear fit is applied to retrieve the evolution rate of q, dq/dt. The rate obtained from simulation and the measured one are then:

$$\begin{aligned} \frac{dq_y}{dt}_{\text{meas}} &= (-0.03 \pm 0.01)/\,\text{min} \\ \frac{dq_y}{dt}_{\text{sim}} &= (-0.04 \pm 0.01)/\,\text{min}\;. \\ \frac{dq_x}{dt}_{\text{sim}} &= (0.00 \pm 0.01)/\,\text{min} \end{aligned} \tag{13}$$

In the horizontal plane, where noise is not present, the q-parameter remains constant over the whole simulation. The transverse distribution in the vertical plane is instead becoming more and more Gaussian as time passes. The good agreement in the vertical plane between the experimental results and the numerical model shows an unprecedented level of control in this kind of beam-profile simulation, together with the high accuracy of the BSRT measurement system, which is capable of capturing such subtle effects with high precision. The “Gaussian shaping” happens with a rate such that, starting from bunches of $q = 1.3$ one would get to a Gaussian in about 10 min, as shown in Fig. 10. It can be then confirmed also in simulations that amplitude noise in the CC acts on the transverse distribution as an equivalent dipolar noise concerning the tail evolution, despite its more complicated nature that links longitudinal and transverse motion.

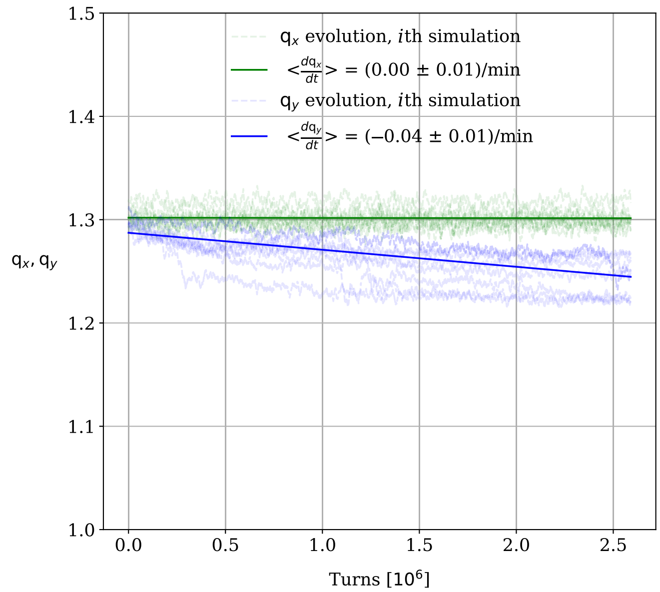

FIG. 11. Evolution of the q-parameter under the effect of amplitude noise in the real SPS lattice, averaged over 10 simulations. Notice that q_x stays constant while q_y is subject to a reduction.

V. CONCLUSIONS AND FUTURE DEVELOPMENTS

In 2023, a new campaign of measurements of emittance growth induced by amplitude noise was performed. The BSRT monitor allowed an unprecedented precision to measure the CERN SPS beam profile. A scan in the octupole strength of the SPS lattice was performed, revealing a dependence of the emittance growth on the octupoles setting. This was not expected by the initially available models. A new simulation was developed with a thin lens model in Xsuite. This allowed reproducing the newest experimental results on the amplitude noise while also reproducing the phase noise simulations and experimental results in the SPS. After carefully looking for the leading causes of this discrepancy, the solution is found to be the linear longitudinal mapping used in the PyHEADTAIL previous simulations. The paper then moved to a deeper level of analysis by looking at how the tail population evolves under the effect of the noise. A q-Gaussian fit is applied to the data as a function of time, showing a decrease in the q-parameter following the effect of noise. This method of analysis is then applied to the SPS lattice with the ability to reproduce the tail population evolution that was measured during the experiment. Together with the emittance growth, this is a demonstration of the unprecedented capability of the simulation setup to reproduce the experimental conditions of the measurements performed in the SPS for the CCs. The study presented here can serve as a foundation for extensive HL-LHC simulation studies regarding q-parameter evolution under the effect of both beam-beam effects and CC noise. Furthermore, in 2025 a horizontal CC has been installed in SPS and a new campaign of measurements has begun. These newly developed simulation tools will guide the analysis and studies that will follow.

ACKNOWLEDGMENTS

This work is supported by the HL-LHC project and The University of Manchester.

DATA AVAILABILITY

The data that support the findings of this article are openly available [41].

APPENDIX: BENCHMARK WITH PHASE NOISE

The reliability of the simulations presented in this study is assessed by benchmarking the new model with both the

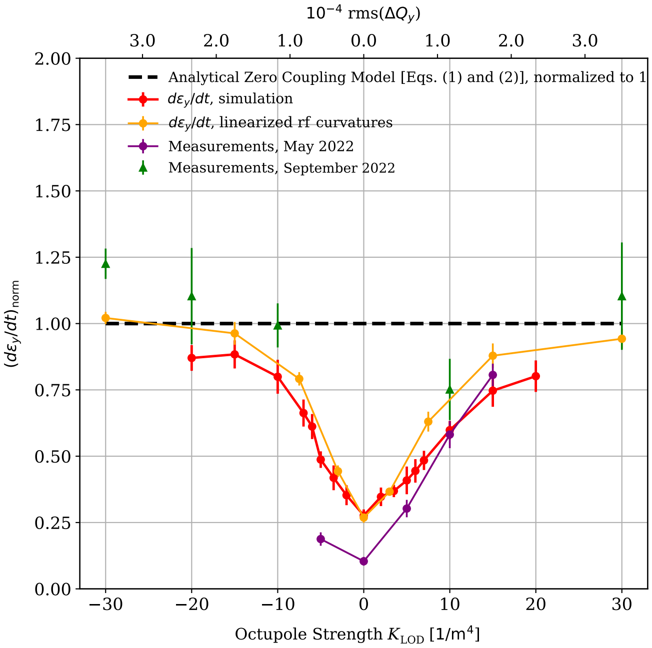

FIG. 12. Measured (magenta: May 2022 and green: September 2022) and simulated (orange) vertical emittance growth driven by phase noise injected in the CC rf system for different octupole settings. Both the nominal simulation (red) and the one with linearized rf curvature (yellow) agree with the data and the previous simulations.

data and the computational studies presented in [13]. Each point in the simulation is the result of $\mathcal{O}(10)$ simulations where the average $\dot{\epsilon}$ is extracted. The results of the benchmark are reported in Fig. 12. First, the new model is capable of reproducing the previous results, with the octupolar dependence found in both the 2022 measurements and the previous PyHEADTAIL model also present here. Second, linearization of the rf curvature leads to the same dependence, showing that in the case of phase noise $\dot{\epsilon}$ variations are not related to the longitudinal map.

[1] O. S. Brüning *et al.*, *LHC Design Report*, CERN Yellow Reports: Monographs (CERN, Geneva, 2004), 10.5170/CERN-2004-003-V-1.
[2] O. Brüning and L. Rossi, The High Luminosity Large Hadron Collider: New machine for illuminating the mysteries of the universe, in *Advanced Series on Directions in High Energy Physics*, edited by O. Brüning (World Scientific, Singapore, 2024), 10.1142/13487.
[3] O. Aberle *et al.*, *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report*, CERN Yellow Reports: Monographs (CERN, Geneva, 2020), 10.23731/CYRM-2020-0010.
[4] W. Herr, Beam-beam interactions, in *Proceedings of the CERN Accelerator School: Intermediate Accelerator Physics, Zeuthen, Germany* (2006), 10.5170/CERN-2006-002.379.
[5] R. Calaga *et al.*, First demonstration of the use of crab cavities on hadron beams, Phys. Rev. Accel. Beams **24**, 062001 (2021).
[6] G. Aad *et al.* (ATLAS Collaboration), The ATLAS experiment at the CERN Large Hadron Collider, J. Instrum. **3**, S08003 (2008).
[7] S. Chatrchyan *et al.* (CMS Collaboration), The CMS experiment at the CERN LHC, J. Instrum. **3**, S08004 (2008).
[8] X. Buffat *et al.*, HL-LHC experiment data quality working group summary report, Technical Report No. CERN-BE-2022-001, CERN, Geneva, 2022, https://cds.cern.ch/record/2802720.
[9] P. Baudrenghien and T. Mastoridis, LLRF studies for HL-LHC crab cavities, in *Proceedings of the 61st ICFA ABDW on High-Intensity and High-Brightness Hadron Beams* (JACoW, Geneva, Switzerland, 2018), THP2WC02, 10.18429/JACoW-HB2018-THP2WC02.
[10] V. Lebedev *et al.*, Emittance growth due to noise and its suppression with the feedback system in Large Hadron Colliders, Part. Accel. **44**, 147 (1994), https://cds.cern.ch/record/248620/files/p147.pdf.
[11] I. Efthymiopoulos, Luminosity estimates including noise from crab cavities (without feedback), in *Proceedings of the Special Joint HiLumi WP2/WP4/WP13 Meeting, Geneva, Switzerland* (2021), https://indico.cern.ch/event/1044711/contributions/4389268/attachments/2264203/3844208/ie_WP2CCemitBU_14.06.2021.pdf.
[12] P. Baudrenghien and T. Mastoridis, Transverse emittance growth due to rf noise in the high-luminosity LHC crab cavities, Phys. Rev. ST Accel. Beams **18**, 101001 (2015).
[13] N. Triantafyllou, F. Antoniou, H. Bartosik, P. Baudrenghien, X. Buffat, R. Calaga, Y. Papaphilippou, T. Mastoridis, and A. Wolski, Impact of beam coupling impedance on crab cavity noise induced emittance growth, Phys. Rev. Accel. Beams **27**, 071001 (2024).
[14] X. Buffat, Suppression of emittance growth by a collective force: Van Kampen approach, in *Proceedings of the IPAC-2022* (JACoW, Geneva, Switzerland, 2022), p. 2197, 10.18429/JACoW-IPAC2022-WEPOTK059.
[15] K. Li, A. Oeftiger *et al.*, PyHEADTAIL, https://github.com/PyCOMPLETE/.
[16] K. S. B. Li *et al.*, Code development for collective effects, in *Proceedings of the ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Hadron Beams (HB-2016), Malmö, Sweden* (JACoW, Geneva, Switzerland, 2016), pp. 362–367, ISBN 978-3-95450-185-4.
[17] Official optics repository of the CERN Super Proton Synchrotron, https://gitlab.cern.ch/acc-models/acc-models-sps.
[18] H. Bartosik, M. Carlà, and K. Cornelis, Reconfiguration of SPS Landau octupole circuits to minimise second order chromaticity, in *Proceedings of 10th International Particle Accelerator Conference (IPAC'19), Melbourne, Australia* (JACoW, Geneva, 2019), MOPGW072, 10.18429/JACoW-IPAC2019-MOPGW072.
[19] MAD-X, https://madx.web.cern.ch/.
[20] IEEE, 1139-2008-IEEE standard definitions of physical quantities for fundamental frequency and time metrology

—random instabilities, 10.1109/IEEESTD.2008.4797525 (2009).

[21] F. Roncarolo, Accuracy of the transverse emittance measurements of the CERN Large Hadron Collider, https://cds.cern.ch/record/1481835 (2005).

[22] R. V. Vene , Overview of beam intensity issues and mitigations in the CERN-SPS fast wire scanners, in *Proceedings of the 15th International Particle Accelerator Conference No. 15, IPAC-2024* (JACoW, Geneva, Switzerland, 2024), pp. 2248–2251, 10.18429/JACoW-IPAC2024-WEPG26.

[23] G. Trad, Development and optimisation of the SPS and LHC beam diagnostics based on synchrotron radiation monitors, MAD-X, https://cds.cern.ch/record/2266055 (2014).

[24] G. Iadarola *et al.*, Xsuite, https://xsuite.readthedocs.io.

[25] G. Iadarola *et al.*, Xsuite: An integrated beam physics simulation framework, in *Proceedings of the 68th Advanced Beam Dynamics Workshop High-Intensity High-Brightness Hadron Beams (HB'23), Geneva, Switzerland* (2023), pp. 73–80, 10.18429/JACoW-HB2023-TUA2I1.

[26] G. Iadarola, Xsuite: An integrated beam physics simulation framework, in *Proceedings of the 15th International Particle Accelerator Conference No. 15, IPAC-2024* (JACoW, Geneva, Switzerland, 2024), pp. 2623–2626, 10.18429/JACoW-IPAC2024-WEPR56.

[27] C. Zannini, SPS impedance model repository, https://gitlab.cern.ch/IRIS/SPS_IW_Model.

[28] L. Palumbo *et al.*, Wake fields and impedance, in *Proceedings of CAS CERN Accelerator School CERN* (1994), pp. 331–390, 10.5170/CERN-1995-006.331.

[29] CERN, Batch system, https://batchdocs.web.cern.ch/index.html.

[30] HTCondor, https://research.cs.wisc.edu/htcondor/.

[31] J. Laclare, Bunched beam coherent instabilities, in *CAS-CERN Accelerator School-Advanced Accelerator Physics*, edited by S. Turner (CERN, Geneva, Switzerland, 1987), pp. 264–326.

[32] S. Furuseth and X. Buffat, Vlasov description of the beam response to noise in the presence of wakefields in high-energy synchrotrons: Beam transfer function, diffusion, and loss of Landau damping, Eur. Phys. J. Plus **137**, 2190 (2022).

[33] N. Triantafyllou *et al.*, Suppression of crab cavity noise induced emittance growth by transverse beam coupling impedance, in *Proceedings of the IPAC-2022* (JACoW, Geneva, Switzerland, 2022), p. 1659, 10.18429/JACoW-IPAC2022-WEOZSP2.

[34] E. M. F. Curado and C. Tsallis, Generalized statistical mechanics: Connection with thermodynamics, J. Phys. A **25**, 1019 (1992).

[35] S. Papadopoulou *et al.*, Impact of non-Gaussian beam profiles in the performance of hadron colliders, Phys. Rev. Accel. Beams **23**, 101004 (2020).

[36] F. Asvesta *et al.*, Characterization of transverse profiles along the LHC injector chain at CERN, in *Proceedings of the IPAC-2013* (JACoW, Geneva, Switzerland, 2023), WEPL158, 10.18429/JACoW-IPAC2023-WEPL158.

[37] K. Umeno and A.-H. Sato, Chaotic method for generating q-Gaussian random variables, arXiv:1205.1690.

[38] T. Sen and J. A. Ellison, HERA-B and halo control using noise, Technical Report No. DESY-HERA-96-09, 1996, https://bib-pubdb1.desy.de/record/172114/files/PUBDB-2014-03565.pdf.

[39] G. Stancari, Diffusion model for the time evolution of particle loss rates in collimator scans: A method for measuring stochastic transverse beam dynamics in circular accelerators, arXiv:1108.5010.

[40] A. V. Igorevic, *Mathematical Methods of Classical Mechanics* (Springer, New York, NY, 1989), 10.1007/978-1-4757-2063-1.

[41] A. Fornara, CC_noise_2023_data: Measurement and simulation data for “Beam profile evolution induced by rf crab cavity amplitude noise in the CERN SPS”, GitHub repository (2025), https://github.com/andreafornara/CC_noise_2023_data.