



Chaos indicators for nonlinear dynamics in circular particle accelerators

C. E. Montanari^{1,4}, R. B. Appleby¹, A. Bazzani^{2,3}, A. Fornara^{1,4}, M. Giovannozzi^{4,a} , S. Redaelli⁴, G. Sterbini⁴, G. Turchetti²

¹ Department of Physics and Astronomy, University of Manchester, Oxford Rd, Manchester M13 9PL, UK

² Dipartimento di Fisica e Astronomia, Università di Bologna, via Iriero 46, 40126 Bologna, Italy

³ Sezione di Bologna, INFN, via Bertoni 6/2, 40127 Bologna, Italy

⁴ Beams Department, CERN, Esplanade des Particules 1, 1211 Meyrin, Switzerland

Received: 17 April 2025 / Accepted: 10 June 2025

© The Author(s) 2025

Abstract The understanding of nonlinear effects in circular storage rings and colliders based on superconducting magnets is a key issue for the luminosity and the beam lifetime optimisation. A detailed analysis of the multidimensional phase space requires a large computing effort when many variants of the magnetic lattice, representing the realisation of magnetic errors or configurations for performance optimisation, have to be considered. Dynamic indicators for chaos detection have proven to be very effective in finding and distinguishing the weakly chaotic regions of phase space where diffusion takes place and regions that remain stable over time scales in the order of multiple hours of continuous operation. This paper explores the use of advanced chaos indicators, including the fast Lyapunov indicator with Birkhoff weights and the reverse error method, in realistic lattice models for the CERN Large Hadron Collider. Their convergence, predictive power, and potential to define a magnetic lattice quality factor linked to long-term dynamic aperture are assessed. The results demonstrate the efficiency of these indicators in identifying chaotic dynamics, offering valuable insights of these chaos indicators for optimising accelerator lattices with reduced computational cost compared to the classical approach based on CPU-demanding long-term tracking campaigns.

1 Introduction

Assessing the chaotic nature of the orbits of dynamical systems is a prominent research field from both the theoretical and experimental points of view. Numerous indicators have been proposed to improve chaos detection efficiency through numerical simulations. These indicators assess the system's response to initial small displacements or random perturbations along a given orbit, providing insights into chaos and long-term stability in the phase space. Recent advances in chaos theory, as well as improved computational resources, have further sparked interest in the potential applications of newer, more refined tools. Dynamic indicators have become essential tools in many fields, including celestial mechanics [1–4] and accelerator physics [5–8], where they help determine the stability of particle motion over extended timescales. For applications in circular accelerators, predicting chaotic behaviour due to multipolar magnetic field errors is especially valuable due to the possibility of studying the diffusion in the phase space and reducing the computational cost of estimating the dynamic aperture (DA). This parameter, widely used to optimise accelerator design, is defined as the extent of the phase-space region where the orbits remain bounded for finite times and is typically computed through long-term tracking campaigns [9–11]. For large and complex accelerator lattices like those of the CERN Large Hadron Collider (LHC) [12], and its planned upgrade, the high-luminosity LHC (HL-LHC) [13, 14], direct particle tracking to study long-term dynamics over physical timescales, such as 1×10^8 turns, equivalent to approximately 2.5 hours of operation, is computationally unfeasible. As a result, there is considerable interest in novel methods and techniques capable of accurately forecasting the long-term stability and evolution of the DA, as well as enabling a greater parametric inspection due to faster DA evaluation.

Chaos indicators have been considered for applications to problems in nonlinear dynamics by various authors, and a brief list of indicators includes: the fast Lyapunov indicator [3, 4] (FLI), its improved version, the Birkhoff weighted fast Lyapunov indicator [15, 16] (FLI^{WB}), the reverse error method [2, 17] (REM), and the frequency map analysis [18, 19] (FMA). Although the first two indicators have yet to be integrated into well-established optimisation frameworks commonly used in accelerator design, the third has seen rapid growth in interest, finding numerous applications within the field of accelerator physics (see e.g. [18–34] for a selected list of references). This gap underscores the largely untapped potential of chaos indicators in practical optimisation tasks in the field of accelerator physics.

^a e-mail: massimo.giovannozzi@cern.ch (corresponding author)

In a previous study [35], we conducted a comprehensive performance evaluation of established and recently developed chaos indicators, focusing on their predictive power for the nonlinear dynamics of a Hénon-like symplectic cubic polynomial map. The FLI^{WB} and the REM indicators turned out to be the most effective in the classification of chaotic orbits with a limited number of iterations, so that we selected them to provide a quantitative analysis of the phase space associated with a realistic LHC lattice. We have also considered the FMA as one of the most popular methods for representing phase space of 4D symplectic maps, and we compared it with the information provided by the chaos indicators.

Yet, applying these results to real-world accelerator structures such as the HL–LHC confronts the complexities inherent in these models, necessitating the optimisation of the algorithm’s computational expenses and the identification of appropriate observables to evaluate the quality factors of a magnetic lattice.

As computational tools continue to evolve, so does the potential to integrate these modern chaos indicators into accelerator physics studies. This study, therefore, represents an initial step in assessing a range of dynamic indicators for accelerator-related models, with the goal of identifying robust methods for understanding and managing chaotic behaviour in particle accelerators, and potentially opening pathways for incorporating these tools into accelerator optimisation frameworks. We aim to demonstrate, employing HL–LHC as an example of a complex and realistic dynamical system, how chaos indicators can offer valuable insight into studying intricate nonlinear beam dynamics. Thus, the application of these advanced tools is undoubtedly broadly beneficial in the field of accelerator physics. However, it is essential to emphasise that the findings of these analyses are notably model-dependent. Consequently, certain observations derived from the HL–LHC scenario may not be universally applicable to other accelerators.

In this paper, we demonstrate how the FLI^{WB} and REM indicators can be applied to realistic HL–LHC lattices to investigate the phase-space structure near the DA. This analysis accounts for the modulational effects introduced by synchrotron dynamics on betatronic motion. We consider extensive numerical simulations of relevant cases of the HL–LHC lattice, analysing the computational cost and the predictability features of the indicators to detect the weakly chaotic regions. For the sake of completeness, we also consider the application of other chaos indicators such as FMA to compare the results with the new indicators that we propose.

We also consider another aspect linked with chaotic regions. In fact, the presence of extended chaotic regions is a necessary condition for diffusive behaviour of the orbits in the phase space, since there is no topological obstruction [36]. We expect that a slow diffusion process occurs in the phase space [37–39] and that a certain fraction of the initial conditions can be lost on a long-term timescale. Therefore, a measure of the Lyapunov exponents can be related to the expected stability time T_s . This quantity can be determined in numerical simulations by defining a control amplitude: the time when an orbit reaches such an amplitude represents the stability time. The details of the system under consideration and the particular physical application determine how the dependence of T_s on the arbitrary value of the control amplitude is managed. This analysis has been performed for polynomial symplectic maps in the neighbourhood of an elliptic point [40–42] suggesting the existence of a power law relation in numerical simulations. We performed a similar analysis for the complicated accelerator lattices with interesting findings.

The paper is organised as follows: Section 2 details how FLI^{WB} and REM chaos indicators are implemented within accelerator tracking codes. Section 3 presents the results of comprehensive numerical analyses, emphasising the prominent features of the chaos indicators examined. Section 4 delves into the potential connection between the information provided by these indicators and the stability issues. Finally, Section 5 summarises the main findings of the investigation. Appendix A offers a brief overview of the chaos indicators, while Appendix B discusses additional numerical simulations conducted with the HL–LHC lattice to mention the features of the chaos indicators in brief.

2 Implementation of chaos indicators in accelerator tracking codes

2.1 GPU parallel tracking of particles

When interactions between individual particles are neglected, such as collective effects or particle-matter interactions involving secondary particle production, the single-particle tracking problem becomes a task that can profit from parallel computing [43]. In fact, this task is suitable for parallelising initial conditions between multiple processors, since each particle’s trajectory can be calculated independently, requiring only shared lattice information. As a result, tracking large numbers of particles can be efficiently parallelised on GPU architectures, which operate under the Single Instruction Multiple Data (SIMD) paradigm [44].

Since 2021, CERN has been actively developing a new symplectic tracking framework known as `Xsuite` [45, 46]. This innovative framework consists of a collection of Python packages designed to enhance the capabilities originally provided by the `SixTrack` code [47]. `Xsuite` boasts the capability of efficient parallelisation on both CPU and GPU architectures, thanks to its ability to generate optimised C code dynamically, starting from a Python implementation. In particular, the `Xtrack` tracking package within `Xsuite` offers the ability to track particles within realistic accelerator lattices on GPU architectures. This advancement enables the tracking of a substantial number of initial conditions in considerably shorter time-frames. The application of GPUs in accelerator physics simulations has already yielded fruitful results in various studies, spanning from investigations into the Hollow Electron Lens to charged-particle tracking studies (see e.g. [48–53]). Among these diverse studies, the GPU version of `Xsuite` played a pivotal role in enabling original statistical analyses of the Hollow Electron Lens [52], and studies of the evolution of the emittance [54, 55].

The structure of the `Xsuite` codebase makes it easy to implement essential components to compute dynamic indicators within the GPU workflow. For example, the shadow-particle method [16, 56], a standard numerical approach used to compute the FLI^{WB} chaos indicator, can be easily incorporated. This method involves tracking a displaced particle to monitor the evolution of its displacement while periodically normalising the particle's position. Importantly, this framework allows researchers to efficiently track many initial conditions. This capacity to handle a substantial volume of initial conditions is a key ingredient for this study, as the combination of extensive scans of the key parameters of the accelerator lattices and dynamic indicators has the potential to provide deeper insights into the phase-space characteristics of realistic accelerator lattices.

However, we stress that the complexity of the magnetic lattice of HL-LHC prevents the possibility of single-particle tracking up to the timescale of hours of operation. Chaos indicators therefore might serve as valuable tools also for predicting the stability, i.e. if the orbit reaches a prescribed amplitude value that is used as a criterion to define the motion unbounded. This means that we are stretching, more or less implicitly, the original goal of dynamical indicators, which is chaos detection, to that of assessing orbit boundedness, which, effectively, corresponds to determining the DA of the lattice.

2.2 Computational effort of considered chaos indicators

In a previous paper [35] we studied the performance and predictability character of the different indicators to detect weakly chaotic phase-space regions for a polynomial, time-dependent symplectic map (4D modulated Hénon map). The FLI^{WB} and REM indicators turned out to be the most effective in the classification of chaotic orbits using only a limited number of iterations. Given that the 4D Hénon map is a good approximation of the dynamics generated by realistic accelerator lattices, we selected the same indicators to perform a quantitative analysis of the phase space of a realistic HL-LHC lattice. Note that, for the sake of completeness, we have also considered the FMA in the appendix.

In realistic applications, one has to evaluate the computational effort required by the dynamic indicators in terms of memory requirements and additional operations. Here, we summarise the key considerations for each indicator using the shadow-particle method.

Obtaining an analytical expression for the tangent map of an accelerator lattice is, in general, a highly complex task, and in the case of intricate magnetic lattices, it may even be impossible. To address this limitation, the shadow-particle method provides a straightforward numerical approach for assessing the chaotic behaviour within a nonlinear magnetic lattice. However, it is known that the accuracy of this method can be influenced by the choice of the initial perturbation amplitude and the frequency of displacement resets [56], an aspect that will be examined in detail in Sect. 2.4.

For FLI^{WB} , the primary computational effort involves tracking the orbits of two particles: the initial condition and the shadow particle. This computational load increases when examining various initial displacements concurrently. In contrast, both REM and FMA do not require tracking shadow particles. However, REM involves both forward and backward tracking, effectively doubling the computational effort similarly to FLI^{WB} , FMA only needs forward tracking to compute the fundamental frequency (i.e. the betatron tune of the orbit) at different time intervals. The memory requirements for FMA may vary depending on the method used to evaluate the tune. For example, using FFT-based algorithms to calculate the fundamental frequency requires storing the complete orbit, while approaches based on the average phase advance algorithm [57, 58] can be implemented without this requirement. We would like to mention that a notable improvement of the average phase advance method consists of combining it with Birkhoff weights [59] as this method offers super-convergence for the case of regular orbits and can be computed in a single forward tracking pass, eliminating the need to store the entire orbit history. Excessive memory requirements inevitably present challenges, making it not convenient to utilise SIMD architectures effectively, due to limitations in GPU onboard memory and the latency associated with frequent data transfers between GPU and host memory.

2.3 Results of preliminary tracking studies with HL-LHC lattices

To compare the performance of the dynamic indicators, we consider a realistic accelerator lattice implementation based on version 1.5 of the HL-LHC layout and optics [60, 61] for the clockwise beam, the so-called Beam 1. The optical configuration is based on the so-called flat optics, in which the horizontal and vertical beta-functions at the interaction points are different. We conducted single-particle tracking without considering beam-beam interactions at top energy, corresponding to 7.0 TeV. This is just a case study taken as an example, and further studies will explore both different configurations and energies, such as round optics and injection energy.

An overview of the machine settings used for the simulation is given in Table 1. Our tracking simulations started at IP3, where the β functions for the two planes are $\beta_x = 113.22$ m and $\beta_y = 225.40$ m. We implemented the lattice using the MAD-X code [62], with the resulting lattice converted to track with the `Xsuite` library.

The magnetic lattice generated by MAD-X includes 60 realisations, often referred to as seeds, of the magnetic field errors associated with the magnets in the ring. These realisations are introduced to consider the impact of the uncertainty on the measured values of the magnetic field errors on the beam dynamics. Hence, the analysis of the tracking results is usually performed by inspecting all realisations and looking at the statistical distribution of the observed features. In our case, from the set of 60 realisations, we selected two representative samples: one exhibiting the largest stable domain (referred to as “best seed” from now on) and one with the

Table 1 Main parameters of the lattice used for the tracking studies

Parameter	Value
Beta-functions at interaction points, β_x^* , β_y^* [cm]	7.5, 18
Tunes, Q_x , Q_y	62.316, 60.321
Chromaticities, ξ_x , ξ_y	15, 15
Current in the arc octupoles [A]	+100
Half crossing angle, $\theta_c/2$ [μ rad]	250

Note that in all these studies the normalised nominal transverse emittances are the same in horizontal and vertical planes and equal $2.5 \mu\text{m}$. The rms bunch length σ_ζ is 7.61cm




Fig. 1 Left plots: representation of the tracking results of a set of initial conditions using the HL–LHC lattice for Beam 1 at 7.0TeV , without beam-beam interaction. The colour encodes the logarithm of the stability time T_s , i.e. the time over which the motion remains bounded, noting that the tracking has been performed up to 1×10^5 turns. The two plots refer to the worst and best seeds, respectively. The initial conditions are distributed over a 100×100 Cartesian grid in the $x_0 - y_0$ plane, with a ROI highlighted in red to mark boundaries for finer sampling used in the analysis of the dynamic indicators. Right plot: orbits in the longitudinal phase space. Initial conditions with varying ζ values are tracked up to 1×10^3 turns for the best seed, illustrating the well-known pendulum-like structure with unstable fixed points around $\zeta \sim \pm 0.35\text{m}$. Black stars indicate the values of ζ at 0, 1, and 2 units of σ_ζ , chosen for investigating the dynamic indicators in their dependence on the longitudinal dynamics

lowest one (the “worst seed”). The stable domain has been estimated using a 100×100 uniform Cartesian grid of initial conditions of the form $(x_0, 0, y_0, 0)$, corresponding to the transverse plane $x_0 - y_0$. It should be noted that linearly normalised coordinates referred to the closed orbit are used, and the unit is the nominal beam sigma, computed using the nominal value of the emittance (see Table 1). The largest connected set of initial conditions that survived up to 1×10^5 turns defines the stable region for the lattice configuration.

Figure 1 (left plots) presents a visualisation of the stability domain for the worst and the best seeds. The initial conditions defined on the Cartesian grid are coloured according to the logarithm of their stability time T_s , which is the time for which the orbit remains bounded, and the yellow area represents the region inside which the orbits are stable up to the maximum number of turns used in the tracking study, i.e. 1×10^5 . The different extents of the stable region are clearly visible together with the difference in the shape of its boundary. The red boundaries represent the region of interest (ROI) that will be used in the rest of the study of chaos indicators.

Figure 1 (right plot) represents an image of the longitudinal phase space computed using the best seed configuration. Various orbits are shown, and those highlighted in red represent special values of the ζ coordinate, corresponding to 0, 1, and $2\sigma_\zeta$, which have been used in the study of chaos indicators to probe the overall impact of longitudinal dynamics on orbit chaos.

2.4 Considerations on the shadow-particle method

The shadow-particle method requires the choice of the amplitude of the initial displacement ϵ_0 of the shadow particle with respect to the reference particle, and the time interval τ between successive renormalisations of the distance between the shadow particle and the reference particle. The choice of these parameters can significantly influence the evaluation of the FLI and FLI^{WB} (see e.g. [16]), since in long-term simulations on complex systems, like HL–LHC, the numerical precision of floating point operations (i.e. 1×10^{-14}) may affect the computation of the Lyapunov indicator if the values of ϵ_0 and τ are too small. In contrast, small values of these parameters are needed to compute the linearisation of the dynamics around the reference particle. In this respect, it should be mentioned that the definition of the displacement and its regular renormalisation should be performed in normalised coordinates, not in standard physical coordinates, to avoid spurious effects in the numerical calculations.

To address the impact of these parameters on the results of numerical simulations, a comparative study has been performed. The results are shown in Fig. 2, where a pair plot of $\log_{10}(\text{FLI}/n)$ evaluated at 1×10^5 turns is presented using the HL–LHC lattice with the worst seed for various settings of ϵ_0 (top row) and τ (bottom row). To distinguish between regular and chaotic orbits, a bi-modal distribution in the histogram of the chaotic indicator is expected [35]. The distribution of the indicator is clearly influenced by the choice of ϵ_0 and τ , with the bi-modal character ultimately disappearing for certain values of the parameters.

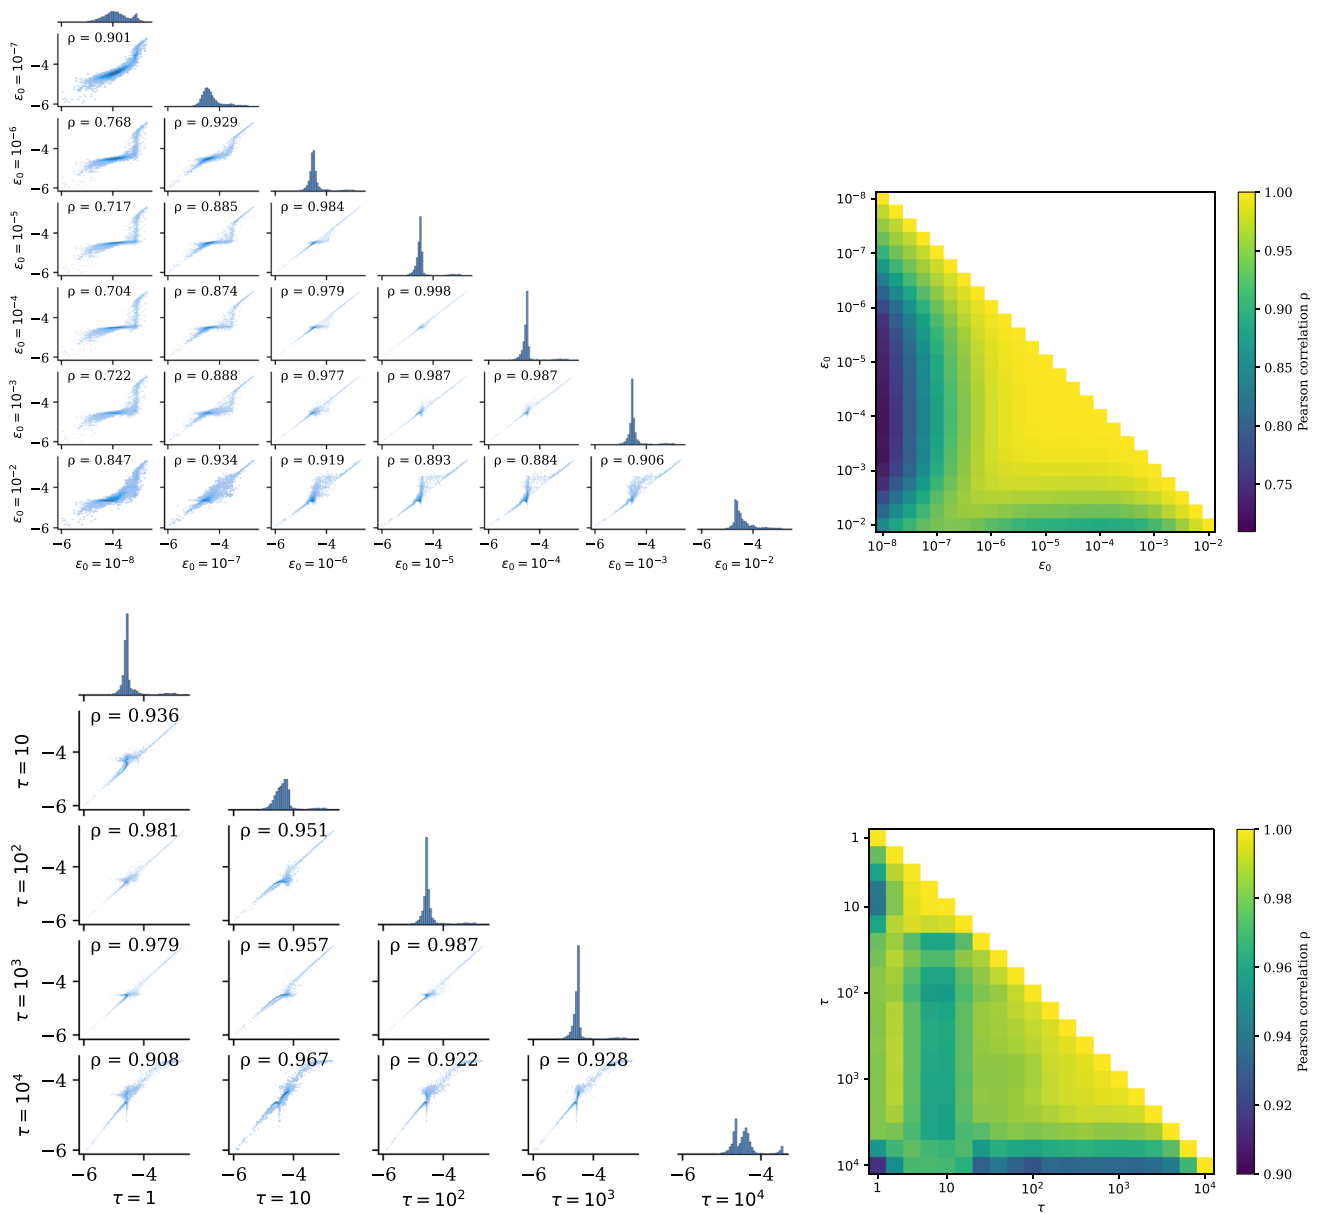


Fig. 2 Pair plots and Pearson correlation analyses of $\log_{10}(\text{FLI}/n)$ evaluated at 1×10^5 turns using the HL-LHC lattice with the worst seed, for various settings for initial displacement ϵ_0 (top row) and the renormalisation interval τ (bottom row). The left column plots show pair plots with histograms of the indicator along the diagonal and correlation plots for different parameter pairings off-diagonally. The right column plots show a finer-grained analysis of the Pearson correlation. For ϵ_0 , lower values (e.g. 1×10^{-8}) yield a wider spread of the indicator for regular initial conditions, similarly to what occurs for large values (e.g. 1×10^{-2}). In contrast, variations in τ are less impactful, provided τ remains below 1×10^3 . For larger τ values, values of $\log_{10}(\text{FLI}/n)$ exceeding -3 appear flattened, pointing to reduced precision at large τ settings

For the remainder of this study, we will consider $\epsilon_0 = 1 \times 10^{-4}$, as it appears to be the best compromise between the two extremes, and the corresponding histogram distribution manifests the sharpest bi-modal distribution.

Regarding the choice of τ , as long as $\tau < 10^3$, we do not observe significant variations in the evaluations of $\log_{10}(\text{FLI}/n)$. Above that value, we have more significant fluctuations in the correlation plots, but the histograms of the indicators do not appear to be significantly affected. This can be understood by the fact that an excessively high value of τ can cause the shadow particle to be too far from the orbit of the reference particle, and its dynamics cannot be used to reliably determine the tangent map of the orbit of the reference particle. Hence, for the remainder of this study, we will consider $\tau = 1 \times 10^2$, as a good trade-off.

3 Chaos-detection studies

3.1 Overview of the chaotic regions of the HL–LHC lattice

A finer Cartesian grid made of 300×300 initial conditions has been considered in the selected ROIs, defined in Fig. 1, to obtain an overview of the phase space that ranges from the neighbour of the closed orbit to the border of the stability region (SR) computed at 1×10^6 turns.¹

The structure of the phase space for the HL–LHC configurations is presented in Fig. 3 and Fig. 4, corresponding to the worst and best seed of Fig. 1, respectively. These figures show a comprehensive view of the behaviour for each configuration. The first row of each figure shows the SR, illustrating the survival of the initial conditions up to 1×10^6 turns for different fixed values of $\zeta = \zeta_0$. In this study, we focus on three values: $\zeta_0 = 0, 1, 2 \times \sigma_z$ with $\sigma_z = 0.0761$ m. The extent of the stable region decreases with an increasing value of ζ_0 for both seeds. It is worth noting that the transition between stable to unstable regions is sharper for the case of the worst seed, whereas the best seed features an almost continuous transition with structures extending to large amplitudes.

To investigate the details of the phase-space structure, we employ chaos indicators. The second row of each figure shows the values of $\log_{10}(\text{FLI}^{\text{WB}}(\hat{x}))$, highlighting the emergence of chaotic regions in the transition between stable and unstable phase space. Moreover, in this scenario, the size of the chaotic regions is influenced by ζ_0 and grows as ζ_0 increases. It is important to note that these indicators identify chaotic regions even within what is considered the stable zone. This observation suggests that there might be initial conditions that, although chaotic, remain stable at least within the number of turns used in our tracking analyses. The exploration of the connection between stable and chaotic dynamics follows in the subsequent section. Additionally, it is noted that while the phase space geometry for the worst-case seed resembles a circle, such is not true for the best-case seed. The chaos indicator effectively points out the intricate structures, and it also exposes islands of regular motion amidst the chaotic regions.

The third row illustrates the phase-space structure derived using the $\log_{10}(\text{REM})$ indicator. In general, the REM results qualitatively align with those shown by FLI^{WB} . It should be noted that the range of values for the FLI^{WB} indicator is quite wide, while that for REM is somewhat narrower, with a distinct clustering around extreme values. This suggests that both REM and FLI^{WB} share a bi-modal distribution of their values. However, REM shows a more pronounced transition between the two modes, which facilitates a method for defining a threshold that distinguishes regular from chaotic orbits.

It should be noted that the chaos indicators have been calculated not just for initial conditions stable for up to 1×10^6 turns, but also for cases with shorter stability times. Specifically, a minimum stability of 1×10^2 turns for FLI^{WB} and 1×10^4 turns for REM has been assumed.

4 Analysis of the beam dynamics using the Lyapunov time

The Lyapunov time T_L , defined as the reciprocal of the maximum Lyapunov exponent (MLE) [63], serves as an indicator of the timescale for detecting chaotic dynamics. The MLE can be estimated directly using the dynamic indicators FLI/n or FLI^{WB} , both of which offer a straightforward calculation of the MLE. Alternatively, the MLE can be derived via REM, as it corresponds to the exponential growth of REM over time in scenarios with chaotic initial conditions.

Within the field of accelerator lattice design and the study of nonlinear beam dynamics, the DA is commonly used as a figure of merit, calculated using numerical simulations performed using the maximum number of turns n_{max} . DA is associated with the stability time T_s , defined as the number of turns required for an orbit to achieve a specified control amplitude. Using fundamental pillars of the field of Hamiltonian dynamical systems, such as KAM theory [64–66] and the Nekhoroshev theorem [67–70], scaling laws with extrapolation capabilities were proposed to describe the long-term evolution of the DA [10, 11, 71], the most recent model being given by

$$D(n) = \rho_* \times \frac{1}{\left[-2e\lambda \mathcal{W}_{-1} \left(-\frac{1}{2e\lambda} \left(\frac{\rho_*}{6} \right)^{1/\kappa} \left(\frac{8}{7}n \right)^{-1/(\lambda\kappa)} \right) \right]^\kappa}, \quad (1)$$

where \mathcal{W} is the Lambert \mathcal{W} function [72, 73]. The variables ρ_* , κ , and potentially λ , unless set to 1/2 following Nekhoroshev original analytical approximation, are free parameters. Here, ρ_* is the apparent radius of convergence for the asymptotic series that converts the nonlinear Hamiltonian system into an integrable Normal Form, predominantly linked to the strength of the system's nonlinearity intensity. κ is linked to the number of degrees of freedom, or the phase-space dimension. The free parameters are determined by fitting the DA model to the results of numerical simulations.

The endeavour to link T_L with T_s continues to pose a challenge investigated in diverse branches of physics, including mathematical physics and astrophysics (see e.g. [40] and related works). T_L is calculated from the MLE and signifies the duration over which an orbit loses the memory of its initial state, with decorrelation primarily in the angle variables. In contrast, T_s results from a computation reminiscent of a first-passage problem, i.e. the direct evaluation of when the orbit of a given initial condition reaches the absorbing boundary condition and is related to the evolution of the action variables. Consequently, T_s holds information about a

¹ Note that this time interval corresponds to approximately 90 s in the actual accelerator.




Fig. 3 Overview of the phase-space structure for the HL-LHC configuration described in the text with the worst seed and tracking performed up to 1×10^6 turns. The three columns refer to the three values of the longitudinal variable ζ_0 used in the tracking studies. The top row shows the stability plots, representing the logarithm of T_s . The middle row presents the behaviour of $\log_{10}(\text{FLI}^{\text{WB}}(\hat{x}))$, highlighting the chaotic regions near the boundaries of the stable area. The bottom row displays the values of $\log_{10}(\text{REM})$, further identifying the transition between regular and chaotic dynamics. The reduction of the extent of isolated islands of regular motion in between of chaotic regions as well as the increase of the chaotic region for increasing values of ζ is clearly visible. Both chaos indicators show the emergence of isolated islands of regular motion in between of chaotic regions. Note that the indicators are computed for stable initial conditions using the tracking data up to 1×10^6 turns, while for unstable initial conditions the computations are carried out using the stable part of the orbit

diffusion coefficient that controls the dynamics of these actions, and this approach has been successfully and effectively advanced in recent research [74–78]. It is conjectured that the diffusive process occurs in phase-space regions where MLE is positive, and hence T_L is finite, and generally $T_L \ll T_s$. This illustrates the importance of analysing the SR identified by the Lyapunov and stability times and verifying if the Nekhoroshev-like scaling law applies to both domains. The DA calculation, which uses T_s , is constrained by the bound $T_s \leq n_{\text{max}}$, with n_{max} being the maximum spins in the simulation. Thus, for initial conditions where $T_s = n_{\text{max}}$, the true stability duration remains uncertain, only allowing the conclusion that $T_s \geq n_{\text{max}}$, which means that the known information provides just a lower bound for the actual stability time. Establishing a direct connection between the Lyapunov time and the stability time could enable the definition of a non-trivial stability time for stable particles from tracking simulations using the Lyapunov time. It should be emphasised that a general relationship cannot typically be established, as evidenced by numerous cases in the literature, such as stable chaos [79, 80] or intermittency [81]. However, even a statistical relationship would be extremely beneficial.




Fig. 4 Overview of the phase-space structure for the HL-LHC configuration described in the text with the best seed and tracking performed up to 1×10^6 turns. The three columns refer to the three values of the longitudinal variable ζ_0 used in the tracking studies. The top row shows the stability plots, representing the logarithm of T_s . The middle row presents the behaviour of $\log_{10}(\text{FLI}^{\text{WB}}(\hat{x}))$, highlighting the chaotic regions primarily at the periphery of the stable area, although some light areas are also present closer to the origin. The bottom row presents the values of $\log_{10}(\text{REM})$, confirming the boundaries between regular and chaotic behaviour. The results demonstrate a more sharply defined stable region compared to the worst seed as well as wider chaotic region. Note that the indicators are computed for stable initial conditions using the tracking data up to 1×10^6 turns, while for unstable initial conditions the computations are carried out using the stable part of the orbit

4.1 Comparison of the properties of the SR based on T_L and of the DA

The determination of the DA is performed using data from the 300×300 Cartesian ROI depicted in Fig. 1. Analysing this grid of initial conditions, the largest connected component of the SR is determined for varying n values, up to a value of 10^7 turns. The area of the SR as a function of n is calculated and used to derive the radius of a circle of equivalent area, representing $D(n)$ as per Eq. (1). The error in this method arises from the grid step size defined for the initial conditions. An analogous method is applied to evaluate similar metrics using T_L : the Lyapunov time is evaluated for the same grid of initial conditions. The largest connected component, bordered by initial conditions with a specific T_L value, is identified, and the radius of the equivalent area circle is calculated. This process is repeated using different Lyapunov time values to define the boundary of the connected region. Figure 5 visually depicts this process, showing the SR related to the stability time (top-left) and the Lyapunov time (top-right), along with the evolution of the DA and the counterpart determined from T_L (bottom). Noticeable differences in the SR geometries corresponding to the two

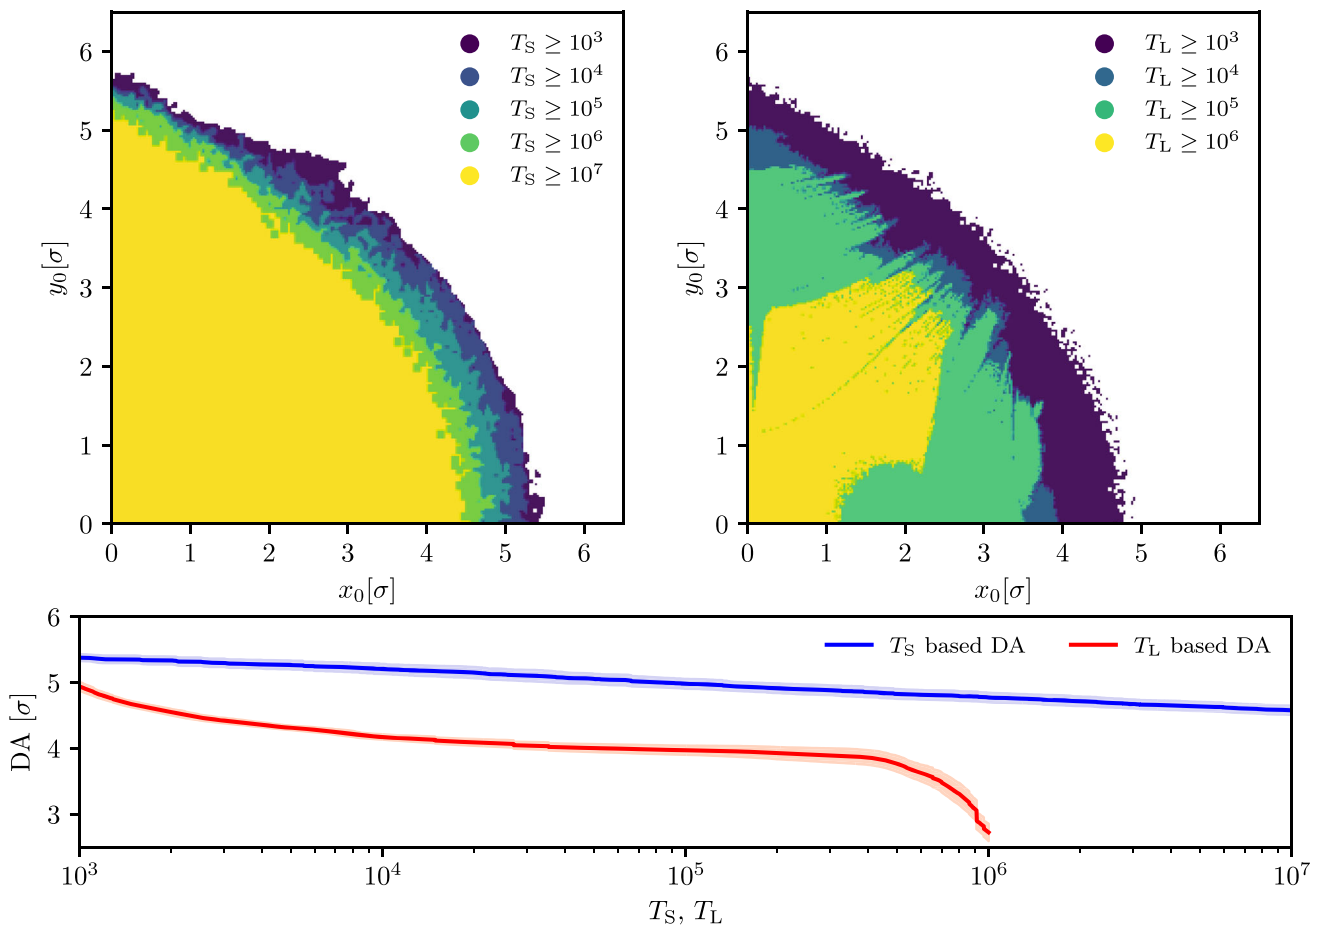


Fig. 5 Top: DA, determined using T_S , (left) and SR, determined using T_L , (right). The boundary of the largest connected component of the stability domain is shown as a function of time. The area of the largest connected component is used to compute the equivalent circular radius, which represents the corresponding DA for a given value of time. The very different geometries of the domains related with T_S or T_L are clearly visible. Bottom: evolution of the DA based on T_S or T_L , a sudden drop in dynamic aperture near $T_L = 1 \times 10^6$ is observed. This feature is due to the irregular erosion of the SR clearly seen in the top-right plot

times are evident. In particular, the geometry becomes highly irregular for the SR tied to Lyapunov time with increasing T_L . This irregularity is reflected in the equivalent radius evolution in the lower plot. Unlike DA, which exhibits a fairly smooth progression, a pronounced decrease is evident as 1×10^6 turns are approached for the case based on Lyapunov time.

Figure 6 provides a detailed comparison of the DA calculated using T_S and T_L for two seeds and three values of ζ_0 . Overall, the standard DA yields smoother curves compared to those based on Lyapunov time. This indicates that the average over the angle in the $x_0 - y_0$ space is more efficient when the analysis is based on the stability time. In particular, when $\zeta_0 \neq 0$, the two curves appear nearly parallel, whereas for $\zeta_0 = 0$, the curves intersect at lower values of the time variable. This pattern becomes even more pronounced when the model is fitted to the numerical data. The construction of the model involves fitting the parameters κ and ρ_* to the observed data. As mentioned previously, κ depends on the dimension of the phase space and is thus expected to be the same for all cases in which longitudinal dynamics is fully included, while a different value is expected when $\zeta_0 = 0$.

The fit parameters obtained are shown in Table 2, together with the calculated reduced chi-squared χ^2_{ν} . The uncertainty in the fit parameters is estimated by observing the variation when the fit is applied to a DA curve adjusted within its uncertainty band, compared to the uncertainty in the fit parameters taken as half of their variation range. For the DA case, the fit is of better quality when $\zeta_0 = 0$, regardless of the seed, and ρ_* appears to be independent of ζ_0 when this variable is different from zero. Furthermore, the fit parameters increase by approximately a factor of two between the case of zero longitudinal variable and nonzero values.

In the case of the T_L -based DA, the quality of the fit is much worse for $\zeta_0 = 0$, while in the other two cases the quality is comparable with that of the fit for the standard DA. These results suggest that when nonlinear contributions are sufficiently strong, the DA computed using T_S and T_L exhibits good agreement. This indicates that DA based on Lyapunov time might serve as a viable alternative tool for analysis in such cases. However, when the system lacks the necessary presence of chaotic structures, this relationship breaks down. In these conditions, the agreement between T_S and T_L deteriorates, as well as the underlying assumptions required to observe the strong evolution of DA over time.

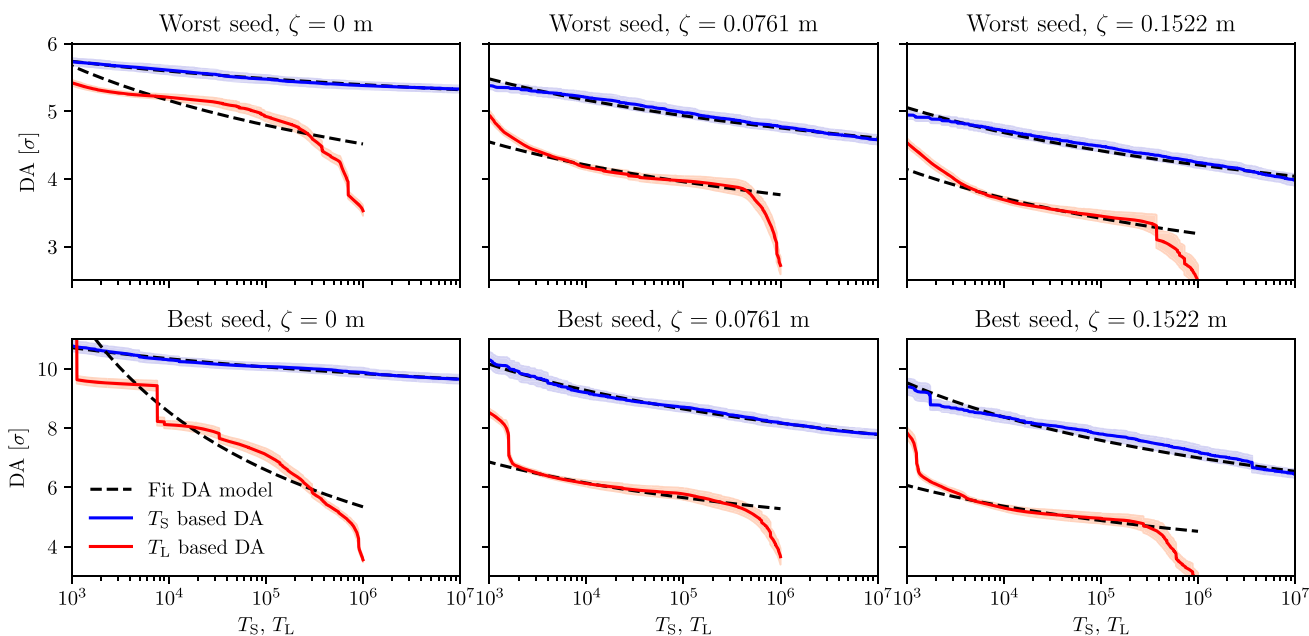


Fig. 6 Comparison between the DA evaluated on T_s (blue lines) and T_L (red lines) evaluated using FLI^{WB} for the two seeds and the three values of ζ_0 considered. A fit of the scaling law (1) is also shown (dotted lines), whose free parameters are reported in Table 2. The agreement between model and numerical data is better for the DA based on T_s , although when T_L is used the agreement with the model improves as a function of ζ_0

Table 2 Fitted parameters for the scaling law (1) applied to the SR determined from T_s and T_L for the cases of the two seeds and three values of ζ_0

Seed	ζ_0 [m]	Fit of T_s based DA			Fit of T_L based DA		
		ρ_*	κ	χ_v^2	ρ_*	κ	χ_v^2
Best	0.0	21.4 ± 0.2	0.121 ± 0.001	0.03	579 ± 3	1.108 ± 0.009	4.7
	0.0761	45.2 ± 1.9	0.31 ± 0.01	0.26	63 ± 9	0.50 ± 0.05	0.29
	0.1522	40.3 ± 1.4			55 ± 8		
Worst	0.0	9.9 ± 0.2	0.091 ± 0.001	0.03	27 ± 1	0.33 ± 0.01	5.01
	0.0761	16.9 ± 0.4	0.217 ± 0.007	0.26	24 ± 3	0.35 ± 0.04	0.21
	0.1522	15.1 ± 0.3			22 ± 3		

The free parameters κ and ρ_* are fitted independently for each individual case with $\zeta_0 = 0$. When $\zeta_0 \neq 0$ the results for the two values of ζ_0 are fitted together to impose the same value of κ . The value of λ is fixed to 1/2 for all fits. The reduced chi-squared χ_v^2 is also reported

4.2 Direct relationship of T_L and T_s

We have also studied the possible direct correlation between T_L and T_s as a function of the initial radial amplitude. The Lyapunov time is an estimate of the exponential loss of correlation between the evolution of two orbits with arbitrarily close initial conditions. This means an exponential increase of the distance of the orbits that is usually due to the phase variables and not action variables, if one considers a weakly chaotic region resulting from the overlapping of nonlinear resonances. The stability time is related to the increase in the amplitude of the orbits up to a fast instability threshold. If the dynamics of actions can be described by a diffusion equation coming from a perturbation, then T_s is inversely proportional to the diffusion coefficient, which is estimated by the square of the strength of the perturbation, whereas if we assume a ballistic increase of the action variables, the stability time is inversely proportional to the perturbation. For weakly chaotic systems, the stability time T_s is mainly related to the action diffusion in the phase space, whereas the Lyapunov exponent is generated by the decorrelation on the angle variables. The relation between T_L and T_s should consider how angle decorrelation can induce diffusive behaviour in action variables and the dependence of stability time on the extension of the chaotic region. This relation is discussed in [40], where the authors consider different regimes to justify a power-law relation $T_s \sim \alpha T_L^\beta$ or an exponential relation $T_s \sim \alpha \exp(\beta T_L)$, where α, β are the free parameters. Their results suggest that there is no universal relationship between T_s and T_L , since the actual regime depends on the specific characteristics of the Hamiltonian system under study.




Fig. 7 Left: stability domain based on T_s together with three circles used to evaluate the average value in different amplitude intervals. Right: mean values of T_s as a function of r_0 for three values of Δr . The best compromise between statistical fluctuations, due to Δr being too small, and loss of information, due to Δr being too large, is achieved for $\Delta r = 0.01\sigma$. These results refer to the case of the worst seed and $\zeta_0 = 0.0761$ m. The same approach is applied to the analysis of T_L

Extensive numerical simulations have been performed to study the correlation between T_s and T_L , as functions of the initial radius $r_0 = \sqrt{x_0^2 + y_0^2}$. We evaluated the averages T_s and T_L for a set of initial conditions in the amplitude interval $[r_0 - \Delta r/2, r_0 + \Delta r/2]$ and in the angular sector comprised between $[0, \pi/2]$. The averages of T_s and T_L are calculated for each interval and the resulting values are plotted as a function of r_0 . We consider T_s computed using numerical simulations performed up to 1×10^7 turns, while T_L is evaluated by FLI^{WB} at 1×10^6 turns.

To find the optimal choice Δr , we first calculate the average values of T_s and T_L for various options for Δr and compare their resulting curves directly. The optimal Δr is the one that strikes the best balance between statistical fluctuations and the preservation of information. Figure 7 illustrates the outcome of the analysis for T_s , with T_L producing similar results, pertaining to the worst seed of the HL–LHC lattice. It is evident that the optimal balance occurs at $\Delta r = 0.01\sigma$, which is the value selected for subsequent analysis. For the uncertainty in the average values of T_s and T_L , we account for the standard deviation of their averages over each interval.

When comparing directly T_L and T_s , as shown in Fig. 8, we can observe that the two curves exhibit a qualitative similar behaviour with respect to the initial amplitude r_0 with a fast increase in stability time up to a saturation value when r_0 is below a certain value. Clearly, the saturation effect is missing for the curve that refers to the Lyapunov time.

We note that average values close to the saturation value of T_s corresponding to 1×10^7 turns are going to be affected by a bias due to the initial conditions that are stable up to the maximum number n_{\max} . The extent of this bias can be observed in Fig. 8, which shows the percentage of samples at different values of r_0 that have not reached n_{\max} . To cope with this problem, only the subset of data corresponding to the white region, i.e. corresponding to the case with $T_s \leq 1 \times 10^7$ turns, has been considered.

In Fig. 9, we show the correlation between the averages, computed at various r_0 values, of T_s and T_L for the two seeds and three values of ζ_0 for the HL–LHC lattice. We observe that the stability time exhibits a clear correlation in a log-log scale plot, regardless of the value of r_0 . To inspect the power-law correlation, which corresponds to the resonance-overlapping regime [40], we fit the data using a power law of the form $T_s = \alpha T_L^\beta$, where α and β are the free parameters. The resulting fit is also presented in Fig. 9 (red dashed lines), together with the values of the fit parameters. We note that, although the power-law fit is computed over a limited range of amplitudes, it agrees well with the complete set of data reported in the plots. The values of the parameters α and β are shown in Fig. 10 as a function of ζ_0 . The parameters for the two seeds are fully compatible, i.e. they are, to a large extent, independent of the seed. A very slight dependence on ζ_0 might be observed.

The power law is capable of capturing the correlation between T_s and T_L with the exponent $\beta \in [2, 4]$, suggesting that the action dynamics could have features compatible with diffusion regimes.

This encouraging outcome indicates that the power law could serve as an approximation of the data, potentially providing extrapolation abilities that merit exploration in accelerator lattice optimisation. Nevertheless, it is important to recognise that these findings can be greatly influenced by the system, suggesting that the power law’s extrapolation abilities need further examination across various scenarios and should be approached carefully.

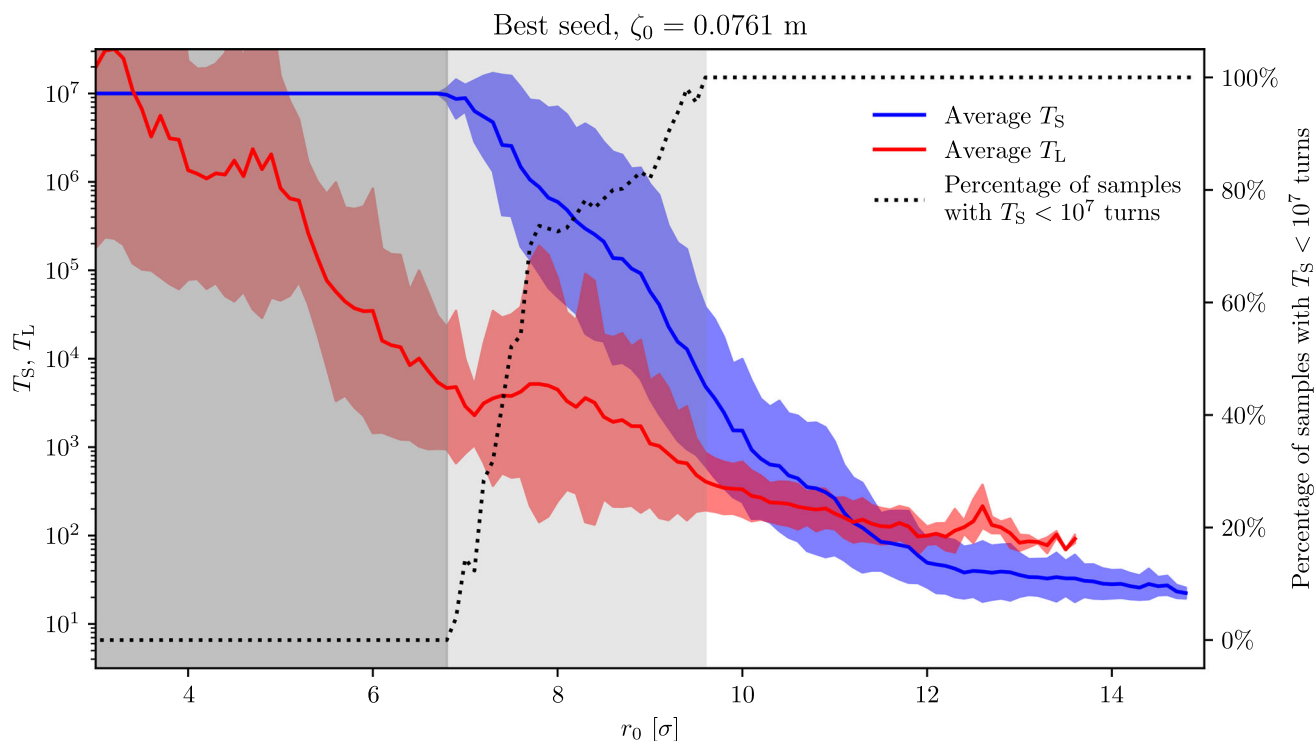


Fig. 8 Average values of T_L and T_S as a function of r_0 for $\Delta r = 0.01\sigma$, for the best seed and $\zeta_0 = 0.0761$ m. T_L is calculated using FLI^{WB} at 1×10^6 turns. The uncertainties on the mean values are represented by the standard deviation of T_S and T_L in each interval. The black dotted line indicates the fraction of samples for which $T_S < 1 \times 10^7$ turns as a function of r_0 . The region shaded in dark grey represents the case for which all samples feature $T_S = 1 \times 10^7$ turns; the region shaded in lighter grey represents the case in which some samples feature $T_S < 1 \times 10^7$ turns; the region shaded in white represents the case in which all samples feature $T_S < 1 \times 10^7$ turns, and this is the region where the fit is performed (see Fig. 9)

5 Conclusions

Chaos indicators serve as an effective means for analysing the phase-space characteristics of a dynamical system, offering insights into chaotic regions and their boundaries while maintaining a low computational cost. This study investigates the application of advanced chaos indicators on realistic HL–LHC lattices to evaluate their usage in accelerator physics and their applicability in optimising accelerator lattices. Specifically, we concentrated on the fast Lyapunov indicator with Birkhoff weights (FLI^{WB}) and the reverse error method (REM), considering their promising results, as demonstrated in previous research on symplectic polynomial maps.

This study compares the performance of the proposed indicators for identifying chaotic behaviour, which is indicated by a positive MLE. This method effectively highlights structures from nonlinear resonances and outlines the boundary of the chaotic layer. It offers a broader application than the FMA method and can be used for realistic accelerator models, like those of the HL–LHC, incorporating synchrotron dynamics. Detecting chaos enables us to pinpoint phase-space regions where action dynamics may become unstable after a certain number of turns, thereby supporting the use of the functional relationship between stability time and the Nekhoroshev optimal estimate for the perturbative series.

We propose that the presence of a positive MLE in a significant portion of phase space validates a diffusion model for the dynamics, enabling the possibility of averaging over fast variables (specifically, the angles) to understand how orbits drift to the boundary of the SR, at least on average. To explore the link between chaos indicators and orbit stability, we analysed the average stability duration alongside the average Lyapunov time at varying amplitudes, employing an angular averaging process akin to that for determining the DA scaling laws. The analysis revealed a distinct power-law relationship between these two times, reminiscent of the resonance-overlap regime outlined in [40]. This indicates that the action dynamics may exhibit a combination of ballistic and diffusion behaviours. Moreover, the power law may offer a reliable extrapolation to higher T_S values, which are critical to optimising accelerator lattices. It is crucial to emphasise that these observations are model-dependent, and other accelerator lattices may produce different results. However, chaos indicators facilitate these analyses in an efficient way.

Inspired by advanced software frameworks that allow efficient GPU parallel tracking of multiple initial conditions, we proceeded with a statistical analysis of the correlation between the conventional SR defined by stability time, i.e. the DA, and the region defined by Lyapunov time. We discovered that both measures follow a similar scaling law, effectively described by a Nekhoroshev-like scaling law. This suggests that both approaches to defining a SR may exhibit analogous long-term behaviour. Consequently, the SR

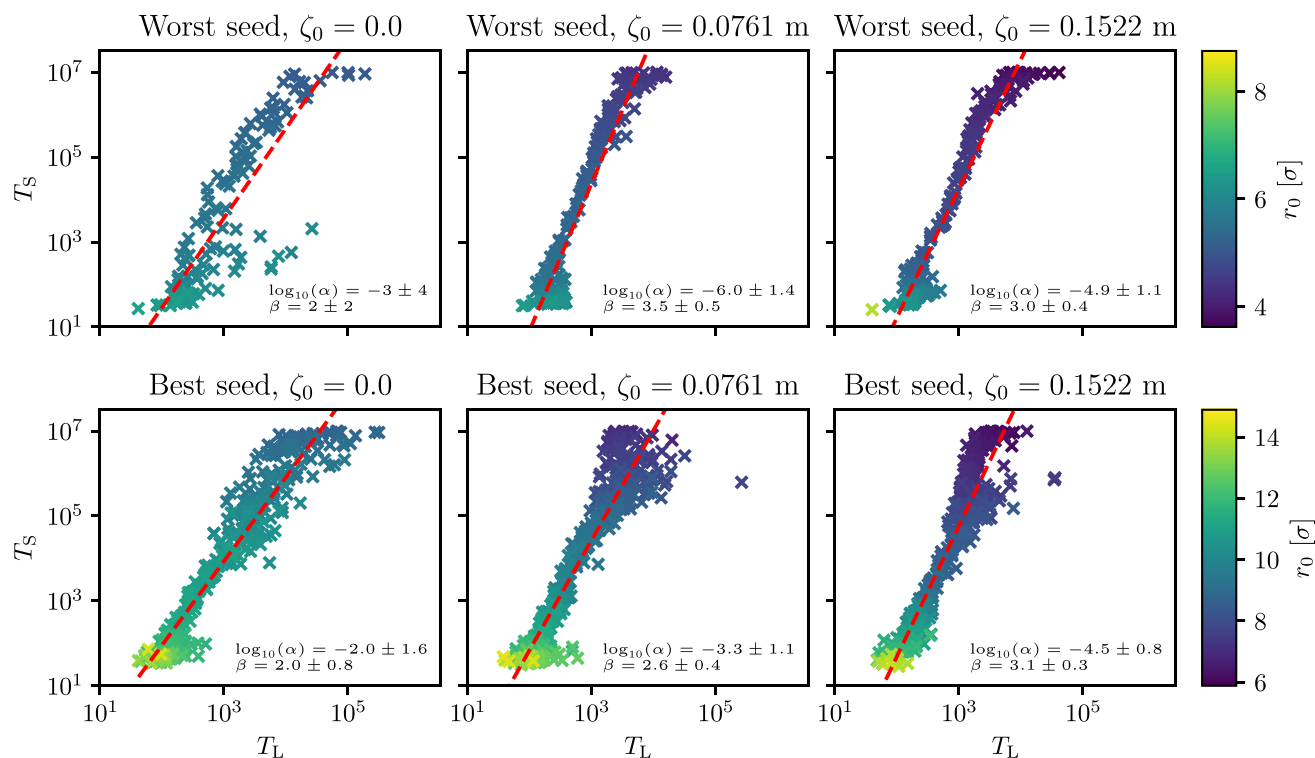
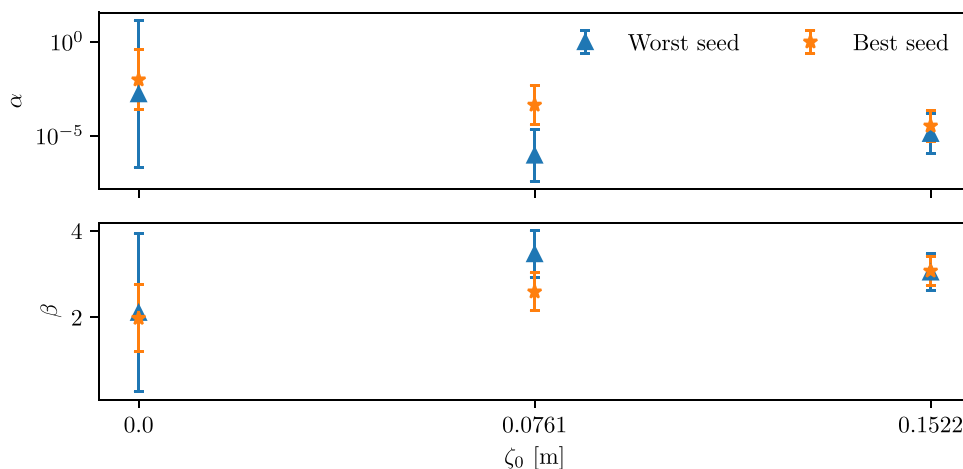


Fig. 9 Correlation plots between T_s and T_L for the two seeds and the three values of ζ_0 considered. T_L is evaluated using FLI at $n = 1 \times 10^6$ turns. The red line represents the best fit of the data in region shaded in white in Fig. 8, using a power-law model $T_s = \alpha T_L^\beta$ typical of the resonance-overlapping regime [40]

Fig. 10 Overview of the fit parameters of the power law between $T_s = \alpha T_L^\beta$ presented in Fig. 9. The parameters for the worst and best seeds are compatible, while the dependence on ζ_0 is clearly visible



based on Lyapunov time might serve as a substitute for the SR based on stability time or the DA. This finding is encouraging because it implies that Lyapunov time can be effectively used to examine the dynamic aperture of an accelerator lattice without executing a comprehensive stability analysis, thereby reducing computational expense.

In conclusion, the encouraging outcomes presented in this study may stimulate a more in-depth investigation into the applicability of chaos indicators within accelerator physics, specifically aimed at optimising accelerator lattices. The objective would be to broaden and implement these methods on realistic accelerator lattices including beam-beam interactions. Exploring the extensive parameter space, comprising elements such as bunch intensities, crossing angle, tunes, chromaticities, amplitude detuning, and β^* , is essential to find the best operational setup, which requires significant computational resources. The efficiency of the chaos indicator could be highly valuable in performing comprehensive analysis of lattice properties while maintaining a feasible computational load.

Funding Open access funding provided by CERN (European Organization for Nuclear Research). Work funded by the HL-LHC project.

Data availability Data sets generated during the current study are available from the corresponding author on reasonable request.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A: Overview of chaos indicators

We recall the fundamental theory and concepts of our selection of dynamic indicators. For a complete description of the individual indicators, we encourage the reader to refer to the sources cited, e.g. [56, 82], and references therein. A similar selection of dynamic indicators was also explored in [35], with an application to a modulated Hénon map with octupolar kicks, for which a full analytical expression of the tangent map is available and tracking up to physically relevant time scales is possible.

Let us consider a non-autonomous Hamiltonian map $M(\mathbf{x}, n)$, defined in \mathbb{R}^{2d} , where d is the number of degrees of freedom and n is the discrete time. We denote by $DM(\mathbf{x}, n)$ the symplectic Jacobian matrix $(DM)_{ij} = \partial M_i / \partial x_j$ evaluated at the point \mathbf{x} at time n .

The orbit of the map \mathbf{x}_n for an initial condition \mathbf{x} , and the recurrence for the tangent map L_n are given by

$$\begin{aligned} \mathbf{x}_n &= M(\mathbf{x}_{n-1}, n-1) & \mathbf{x}_0 &= \mathbf{x} \\ L_n(\mathbf{x}) &= DM(\mathbf{x}_{n-1}, n-1)L_{n-1}(\mathbf{x}) & L_0 &= I. \end{aligned} \tag{2}$$

The chaotic character of an initial condition is thus probed by using the linear response of the system to different combinations of stochastic deviations via dynamic indicators.

A.1: Fast Lyapunov indicator

The fast Lyapunov indicator (FLI) [83] is one of the most established Lyapunov-based dynamic indicators, due to its straightforward implementation and sensitivity to identify chaotic structures, such as the Arnold web [3]. It accomplishes this by providing a rapid numerical estimate of the MLE.

The FLI dynamic indicator considers the linear response of the system when subjected to small stochastic perturbations of the initial condition \mathbf{x} , specifically $\mathbf{y}_0 = \mathbf{x} + \epsilon \boldsymbol{\xi}$, where $\boldsymbol{\xi}$ is a random vector with a zero mean and a unit covariance matrix. The linear response vector, denoted as $\boldsymbol{\Xi}_n(\mathbf{x})$, is defined as:

$$\boldsymbol{\Xi}_n(\mathbf{x}) = \lim_{\epsilon \rightarrow 0} \frac{\mathbf{y}_n - \mathbf{x}_n}{\epsilon} = DM(\mathbf{x}_{n-1}, n-1)\boldsymbol{\Xi}_{n-1} = L_n(\mathbf{x})\boldsymbol{\xi}. \tag{3}$$

The FLI at the n th iteration for a perturbed initial condition \mathbf{x} is given by:

$$FLI_n(\mathbf{x}) = \log \|\boldsymbol{\Xi}_n(\mathbf{x})\|. \tag{4}$$

The ratio $FLI_n(\mathbf{x})/n$ converges to the MLE as n tends to infinity. For initial conditions associated with regular orbits, where the largest Lyapunov exponent is zero, $FLI_n(\mathbf{x})/n$ follows a power-law convergence to zero.

The numerical evaluation of FLI depends on the initial choice of $\boldsymbol{\xi}$, which can lead to the detection of different spurious structures in phase space [35]. To address this problem, one considers the statistical properties of multiple $\boldsymbol{\xi}$ displacements or employs more advanced indicators such as the orthogonal fast Lyapunov indicator (OFLI) [84], which allows one to examine the full spectrum of Lyapunov exponents. Alternatively, one can opt for an invariant definition of the dynamic indicator, such as the Lyapunov Error (LE) [85], which provides a Lyapunov-related dynamic indicator independent of the choice of the initial perturbation.

If an analytic expression for $DM(\mathbf{x}, n)$ is not available, estimating the linear response can be performed by the shadow-particle method [56]. This method is based on the tracking of two nearby initial conditions \mathbf{x}_n and \mathbf{y}_n , evaluating the estimate $\boldsymbol{\Xi}_n(\mathbf{x}) = (\mathbf{y}_n - \mathbf{x}_n)/\epsilon$ where ϵ is the initial distance. This is performed by scaling the displacement vector $\mathbf{y}_n - \mathbf{x}_n$ every τ iterations to the initial length ϵ and computing the product of the scaling factors. Such a procedure prevents spurious effects in the FLI evaluation arising from excessive orbit drifting. For clarity, this periodically modified \mathbf{y} is denoted as \mathbf{y}' , and its formulation is as follows:

$$\begin{aligned} \mathbf{y}'_0 &= \mathbf{y}_0; \\ \mathbf{y}'_i &= DM(\mathbf{x}_{i\tau-1}, i\tau-1) \cdots DM(\mathbf{x}_{(i-1)\tau}, (i-1)\tau) \left[\mathbf{x}_{(i-1)\tau} + \epsilon \frac{\mathbf{y}'_{i-1} - \mathbf{x}_{(i-1)\tau}}{\|\mathbf{y}'_{i-1} - \mathbf{x}_{(i-1)\tau}\|} \right], \quad 0 \leq i \leq n/\tau. \end{aligned} \tag{5}$$

To incorporate y'_n into the FLI expression, we can take advantage of the properties of the logarithm in Eq. (4). This allows us to express $FLI_n(x)/n$ as an average along the trajectory x_n , resulting in:

$$\frac{FLI_n(x)}{n} = \frac{1}{n} \sum_{i=1}^{n/\tau} \log \left\| \frac{y'_i - x_{i\tau}}{\epsilon} \right\|. \tag{6}$$

It should be noted that it has been emphasised [63] how the choice of both the value of ϵ and the time interval τ between displacement resets to ϵ can significantly influence the final evaluation of $\Xi_n(x)$ and, consequently, $FLI_n(x)$. This dependence on the choice of ϵ and τ implies the need to perform a convergence study for each specific application of FLI, to ensure that the chosen values of ϵ and τ are suitable for the problem at hand.

A.2: Birkhoff weighted fast Lyapunov indicator

In the case of the quasi-periodic time series, such that associated with regular orbits in the phase space, the use of non-uniform weights to compute an average quantity can yield superior convergence properties compared to a standard uniform average [16]. In the work of Das et al. [15], it is presented how employing the Weighted Birkhoff averaging method WB_n to evaluate the maximum Lyapunov exponent provides accelerated convergence rates. Furthermore, improvements in FLI evaluation convergence rates were observed in cases involving a modulated Hénon map with octupolar kicks [35].

The Weighted Birkhoff averaging applies a weighting function $w\left(\frac{i}{n}\right)$ to the underlying data, akin to a window function in frequency space. An especially effective weighting function in enhancing the convergence of quasi-periodic time series averages [16] is:

$$w(t) := \begin{cases} \exp\left[-\frac{1}{t(1-t)}\right], & \text{for } t \in (0, 1) \\ 0, & \text{for } t \notin (0, 1) \end{cases}. \tag{7}$$

Substituting the standard mean with $w(t)$ in Eq. (6) we get:

$$FLI_n^{WB}(x) = \frac{1}{\tau} \sum_{i=0}^{n/\tau} w\left(\frac{i}{n/\tau}\right) \log \left\| \frac{y'_i - x_{i\tau}}{\epsilon} \right\|, \tag{8}$$

and in scenarios where the series $\log \left\| \frac{y'_i - x_{i\tau}}{\epsilon} \right\|$ exhibits quasi-periodic behaviour, that is, non-chaotic behaviour, $FLI_n^{WB}(x)$ will converge more rapidly than $FLI_n(x)/n$.

A.3: Reverse error method

The reverse error method (REM) [2, 17, 85, 86] is based on the linear response of a dynamical system (susceptibility) to small stochastic perturbations applied over n iterations of the map M , followed by n iterations of the inverse map M^{-1} . This process can be described as follows:

$$\begin{aligned} y_{n'} &= M(y_{n'-1}, n' - 1) + \epsilon \xi_{n'} & 1 \leq n' \leq n \\ y_{n'} &= M^{-1}(y_{n'-1}, 2n - n') + \epsilon \xi_{n'} & n + 1 \leq n' \leq 2n, \end{aligned} \tag{9}$$

where $y_0 = x$. The linear response at iteration $2n$ is denoted as:

$$\Xi_{Rn}(x) = \lim_{\epsilon \rightarrow 0} \left\langle \frac{y_{2n} - x}{\epsilon} \right\rangle, \tag{10}$$

and the REM dynamic indicator is defined as:

$$\left(\text{REM}_n(x) \right)^2 = \frac{1}{2} \left(\frac{y_{2n} - x}{\epsilon} \right)^2, \tag{11}$$

which represents the Euclidean distance between the initial condition and the resulting displaced particle after the forward and backward tracking in the limit of zero noise amplitude averaged on the noise realisations.

The standard approach in REM treats numerical round-off as a form of pseudo-random deviation along the orbit, with an amplitude typically on the order of $\epsilon \sim 10^{-16}$ in the 8-byte standard IEEE754 [87] representation of real numbers. This round-off is akin to white noise in terms of its spectral characteristics.

The key aspects of REM include its ease of implementation, particularly when an expression for the inverse map is available, and the result is independent of the choice of the initial perturbation. However, since we consider a single realisation of the round-off noise, we may have relevant fluctuations as a function of the iteration number n .

By treating the numerical round-off error as a white-noise process, it is expected to observe a power-law evolution of REM over time for regular orbits, while for chaotic orbits it is expected to observe an exponential growth determined by their MLE. This difference in growth rates can be exploited to achieve a fast binary classification of regular and chaotic orbits.

A.4 Frequency map analysis

The Frequency Map Analysis (FMA) examines how the main frequencies of an orbit evolve over time, leveraging the quasi-periodic nature of regular orbits of Hamiltonian systems (i.e. KAM tori). J. Laskar first introduced the FMA in the field of celestial mechanics, where it quickly spread to other disciplines, yielding significant outcomes, particularly in accelerator physics (see e.g. the following selected references [18–23, 25–34]).

Given a Hamiltonian system of n degrees of freedom in action-angle variables $H(I, \theta) = H_0(I) + \varepsilon H_1(I, \theta)$, where for $\varepsilon = 0$ the Hamiltonian $H_0(I)$ is integrable. If the system is non-degenerate

$$\det\left(\frac{\partial v(I)}{\partial I}\right) = \det\left(\frac{\partial^2 H_0(I)}{\partial I^2}\right) \neq 0, \quad (12)$$

the application

$$\begin{aligned} F : \mathbb{R}^n &\longrightarrow \mathbb{R}^n \\ (I) &\longrightarrow (v) \end{aligned} \quad (13)$$

is a diffeomorphism between the action and the frequency space. This means that the invariant tori are equally identified by the action variables I or by their corresponding frequency vector v . For a non-degenerate system, when ε is sufficiently small, the KAM theorem [64–66], asserts that there still exists a set of initial conditions for which the perturbed system still possesses regular motion, that is, KAM tori set, for which, according to Pöschel's theorem [88], a similar diffeomorphism still exist even if the KAM tori define a Cantor set.

Within this theoretical framework, one distinguishes between regular orbits, which exhibit a discrete Fourier spectrum, namely that of the KAM tori, and chaotic orbits, whose Fourier spectrum is complex since the corresponding orbit does not lie on a torus [27].

FMA consists in performing numerical evaluations of the frequency vector v from a time series corresponding to a certain interval $[i, i + 2n - 1]$, determining the Euclidean distance between the frequencies computed using the data in the time interval $[i, i + n - 1]$ and $[i + n, i + 2n - 1]$, and then inspecting the evolution of the distance as a function of n . In case of a regular orbit, the distance tends to zero, the frequency vector converging to the true values to the precision of the numerical method used to determine the orbit frequency. In the case of a chaotic orbit, the distance will stay away from zero, showing a sort of diffusion in the frequency space [25].

To efficiently estimate numerically the fundamental frequencies of a time series, multiple studies have been performed to improve standard algorithms such as the fast Fourier transform (FFT) or the average phase advance (APA) [18, 58, 89, 90]. In the work of Bartolini et al. [57], the fundamental frequency is evaluated using an FFT combined with a Hanning filter and an interpolation algorithm. In recent studies [59], the APA algorithm is improved by applying the weighted Birkhoff averaging [16], which will be used in this paper to perform the FMA evaluation. In the context of this research, we will use the APA algorithm improved with the weighted Birkhoff averaging to estimate the fundamental frequencies.

To compare FMA with the other dynamic indicators, we define FMA_n as the Euclidean distance between two vectors of fundamental frequencies v_1 and v_2 , evaluated, respectively, over the time intervals $[0, n/2]$ and $[n/2, n]$ for an orbit. An initial condition on a regular torus will exhibit regular motion and have FMA_n converging to zero. In contrast, for an initial condition that is not in a regular torus, the distance between v_1 and v_2 is bounded away from zero.

B: Features of dynamic indicators

B.1: Use of Birkhoff weights with FLI

To quantify the convergence improvements given by the Birkhoff weights, we compare the values obtained for FLI at different times for one of the HL–LHC lattices, using the standard approach that considers the mean in Eq. (6), that is, FLI/n , or the weighted mean based on the use of Birkhoff weights as in Eq. (8), that is, FLI^{WB} .

In this analysis, we consider two ensembles of regular and chaotic particles that have been classified by means of the FLI indicator value computed for $n = 1 \times 10^5$ turns. Inevitably, this evaluation does not reach an absolute level of reliability due to the limited number of turns inspected. Therefore, there could be some misclassified particles in our sampling. To overcome this issue as much as possible, we consider regular particles that have reached a final value of $\log_{10}(FLI/n) < -4.5$, and, as chaotic particles, we consider those that have reached a final value of $\log_{10}(FLI/n) > -4.5$. This arbitrarily threshold strategy is based on the expected properties of the FLI dynamic indicator and considers a subset of initial conditions that have already manifested clear regular or chaotic behaviour at 1×10^5 turns, while excluding those that do not yet have a clear classification.

The sets are then used to inspect the time evolution of both FLI/n and FLI^{WB} , to assess possible classification improvements in the latter compared to the first. These improvements can be, for example, an increased convergence rate in the indicator value or an increased spread between the values of regular and chaotic initial conditions.

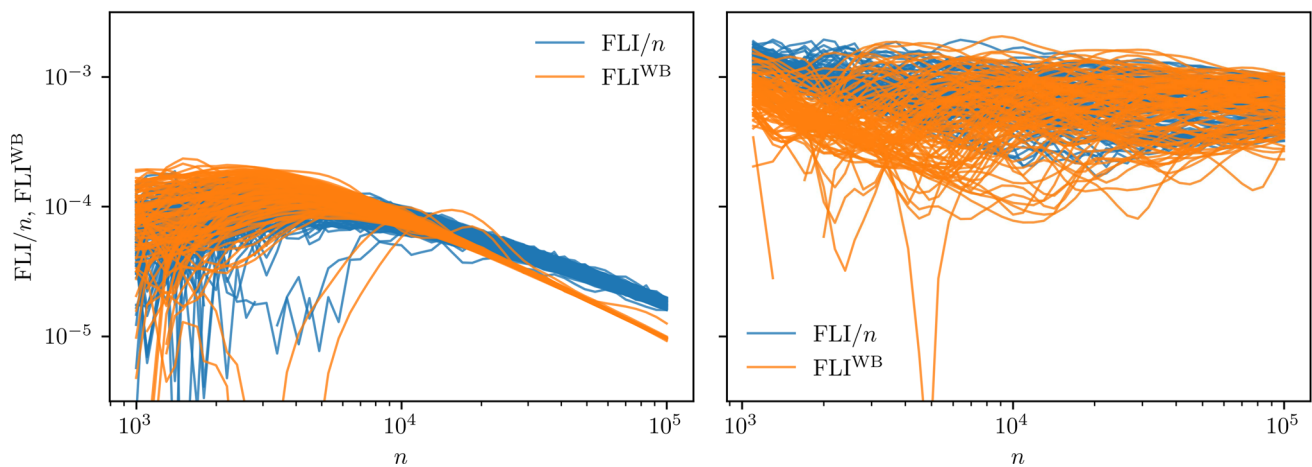



Fig. 11 Time evolution of FLI computed using either a standard mean (FLI/n) or the Birkhoff averaging (FLI^{WB}). Left: indicators computed for a set of 100 regular initial conditions, above 10^4 turns, a slightly faster convergence to zero is observed for FLI^{WB} . Right: indicators computed for a set of 100 chaotic initial conditions. A slight difference in convergence rate is observed for low n values, before reaching a saturation value of the indicator of the order of $10^{-3} - 10^{-3.5}$. The results are obtained using the worst seed and $\zeta_0 = 0.0761$ m

Fig. 12 Distribution of values of the various dynamic indicators as a function of time for a realistic HL-LHC lattice. For low values of the number of turns n , the distribution is in general represented by a unimodal function. For higher values of n , we can see the formation of either two separate clusters, making the distribution bi-modal, or an individual cluster with a significant tail. $\log_{10}(FMA)$ constitutes an exception, as it evolves forming a tri-modal distribution. (HL-LHC lattice used: worst seed, $\zeta_0 = 0.0761$ m.)



In Fig. 11 (left), the comparison between the two indicators is made for a subset of the set of regular initial conditions. It is possible to observe how, for regular initial conditions, Birkhoff averaging does slightly improve the convergence rate, as the indicator tends to the zero value. This improvement can be clearly observed above the sampling of $n = 10^4$ turns, while before that value the two indicators appear to have similar strong fluctuations around the range of values $[-3.75, -5.0]$. This initial behaviour might be related to the low number of turns considered, along with the renormalisation time interval $\tau = 10^2$, which finally causes the number of samples to be less than 100 for both averaging methods, possibly leading to a delay in the appearance of the expected convergence rate.

In Fig. 11 (right), we show the comparison for the subset of chaotic initial conditions. In this case, a saturation region is observed for the indicator value of the order of $10^{-3} - 10^{-3.5}$ for both indicators. When this value is reached, both indicators oscillate around it. However, the slope with which this nonzero value is reached is different for the two indicators and is higher in absolute value




Fig. 13 Colour maps of the various dynamic indicators for a realistic HL–LHC lattice, using the worst seed and $\zeta_0 = 0.0761$ m, evaluated at $n = 1 \times 10^5$ turns. It can be seen how the indicators globally highlight the same structures in phase space, except for FMA, which also shows additional structures

for FLI/n than for FLI^{WB} . This, combined with the fact that both indicators showed a similar convergence rate for regular initial conditions, suggests a greater difference in the convergence rate of regular and chaotic initial conditions for FLI^{WB} , compared to FLI/n . We must also point out how some chaotic initial conditions exhibit large fluctuations in FLI^{WB} before the saturation point, reaching values comparable to the regular initial conditions. These isolated cases might be artefacts caused by the Birkhoff weights, which amplify certain modes in the time series when the chaotic behaviour still has not fully manifested.

Despite these isolated oscillations, the improvement brought about by the Birkhoff averages can be appreciated by comparing the evolution of the value distribution of $\log_{10}(FLI/n)$ and $\log_{10}(FLI^{WB})$ reported in Fig. 12, since there the time evolution of the value distribution is shown for all the initial conditions shown in Fig. 13. The part of the distribution corresponding to the regular initial conditions reaches its peak (yellow band) and moves towards zero with increasing n . However, the displacement towards zero is faster for FLI^{WB} , and the peak is also narrower, potentially allowing for a better binary classification of regular and chaotic orbits, as pointed out in [35].

B.2: Dependence of FMA from the longitudinal dynamics

As we can observe in Fig. 13, the chaotic structures highlighted by FMA are very different from those highlighted by the other dynamic indicators. This is because FMA is generally sensitive to tune changes, which are not necessarily related to chaotic dynamics but rather related to the presence of resonances or tune modulation.

Another feature emerges when considering the evolution of the distribution of indicator values as a function of time. In fact, if we compare the evolution of the value distribution of FMA, we can observe, in Fig. 12, how FMA, unlike the other dynamic indicators, does not show a tendency to create a sharp bi-modal distribution.

To further highlight this characteristic of FMA, we can observe how the indicator is particularly sensitive to the presence of longitudinal dynamics. Indeed, the longitudinal dynamics couples with the transverse one, also introducing tune modulation via a finite value of the chromaticity. In Fig. 14, we present the FMA evaluated for the worst seed for the three values of ζ_0 inspected, and we compare the resulting structures with those highlighted by FLI^{WB} . Two essential features can be observed. The first is that the




Fig. 14 $\log_{10}(\text{FMA})$ (top row) and $\log_{10}(\text{FLI}^{\text{WB}}(\hat{x}))$ (bottom row) both evaluated on the same seed on the worst seed and for the three values of ζ_0 at $n = 1 \times 10^5$ turns. The differences between the two indicators are enhanced for larger values of ζ_0

chaotic regions at the border of the stable region resemble very much the two indicators, which show that their behaviour is similar. The second is a strong difference between the two indicators in the region close to the vertical and horizontal axis. This difference grows significantly as a function of ζ_0 .

With increasing values of ζ_0 , the chaotic region detected by the FMA extends towards the origin along the vertical axis. Furthermore, the width of this chaotic region increases with ζ_0 . It is quite clear that the chaotic behaviour detected by FMA is mainly related to the presence of a strong modulation of the tunes. In this sense, this observation suggests that the use of FMA to identify chaotic regions in phase space is taken with a grain of salt whenever modulation of linear tunes is present: The use of other dynamic indicators could effectively disentangle the artefacts due to tune modulation from genuine chaotic behaviours.

B.3: Overview of REM behaviour

In the performance review presented in [35], REM proved to be the most effective dynamic indicator to quickly achieve a binary distinction between regular and chaotic initial conditions. This is mainly due to the different scaling law followed by the indicator for regular and chaotic initial conditions, which leads to a sharper bi-modal distribution of the indicator values.

When applying REM to the HL–LHC lattices, we observe with high clarity the chaotic regions that are present in the phase space. In Fig. 13, we can qualitatively compare the results obtained by REM, calculated for an example seed, with those obtained with FLI^{WB} . We can observe how REM tends to highlight chaotic regions qualitatively sharper than FLI^{WB} . However, we stress that, differently from previous research based on simpler accelerator models, it is not possible to establish a ground truth for the chaotic regions in the realistic HL–LHC lattice. However, we expect, based on the results on simpler models, that REM is able to better highlight chaotic regions in the phase space.

The behaviour of the distribution of the indicator values can be appreciated by comparing the evolution of the value distribution of $\log_{10}(\text{REM})$ and $\log_{10}(\text{FLI}^{\text{WB}})$ as a function of time, as shown in Fig. 12. In fact, $\log_{10}(\text{REM})$ tends to a bi-modal distribution much faster than $\log_{10}(\text{FLI}^{\text{WB}})$ and in such a way that a threshold for detecting chaotic initial conditions would be almost independent of n .

A more detailed overview of the behaviour of REM for different initial conditions is presented in Fig. 15, where a comparison is presented between the power-law evolution of REM of a regular initial condition and the exponential evolution of a chaotic one.

Fig. 15 Evolution of REM for a regular (blue) and a chaotic (orange) initial condition. The case of the regular initial condition shows a power-law behaviour, while that of the chaotic initial condition shows an exponential-like behaviour with a saturation value of the order of 1×10^1 . An exponential fit of the chaotic initial condition to reconstruct the value of the MLE is also presented

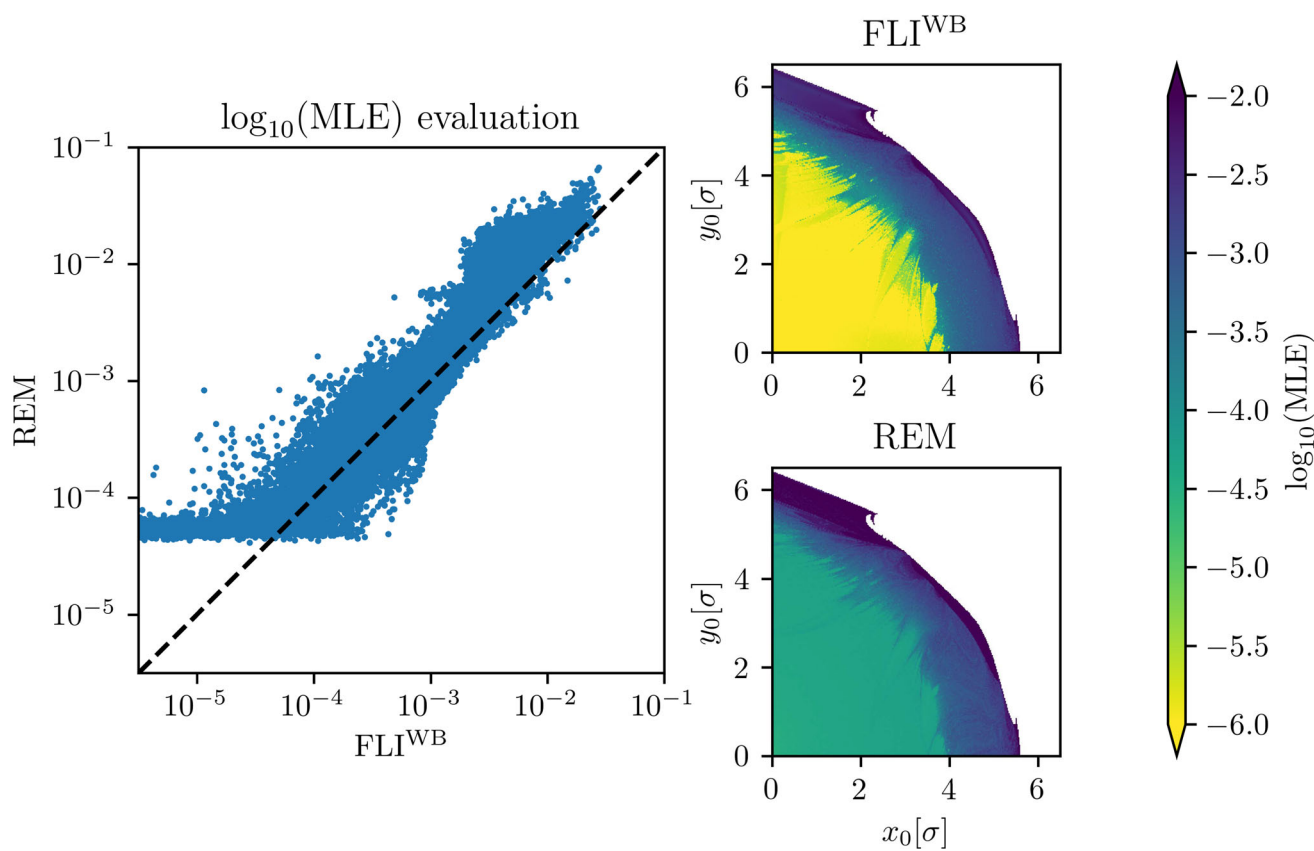
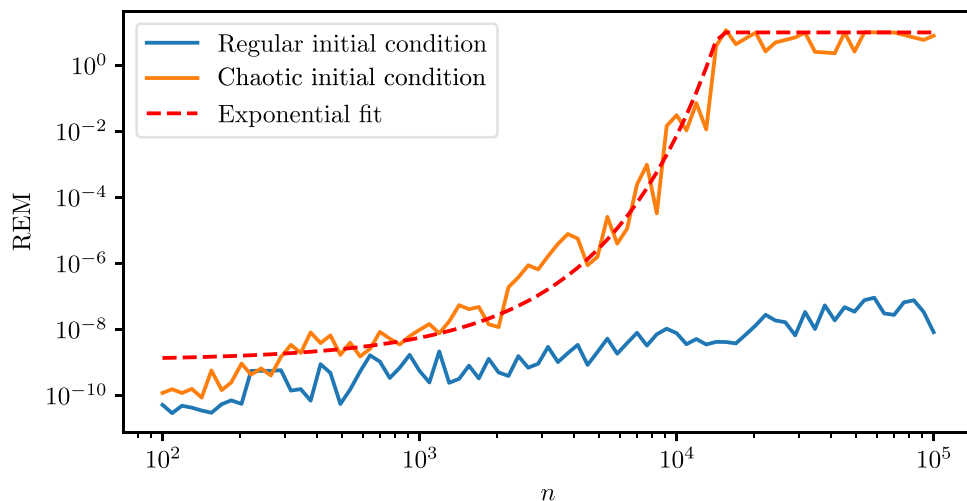


Fig. 16 Right: value of MLE for the Cartesian grid of initial conditions reconstructed using REM and FLI^{WB} for the case of the worst seed. Left: correlation plot of MLE reconstructed using REM and FLI^{WB}

Note how REM, for the chaotic initial condition, saturates to values up to 10^1 . This saturation value is comparable to the diameter of the explorable phase space, and it indicates that the chaotic behaviour eventually led the backtracked particle to completely lose track of the initial path, exploring completely different regions of the phase space.

By fitting an exponential law to the evolution of REM, while also excluding saturated parts of the data, it is possible to recover the value of the MLE governing the timescale of the process. We present the value obtained for all initial conditions and compare it to the value obtained from FLI^{WB} in Fig. 16. We can see how the two indicators agree well, except for a few outliers, which might be related to cases in which the chaotic behaviour has not yet fully manifested for the FLI^{WB} indicator to be able to correctly evaluate the MLE.

The promising performance of REM in highlighting chaotic structures in the phase space, along with its straightforward numerical implementation, makes it a very interesting tool for studying phase-space structures in realistic accelerator lattices.

References

1. F. Panichi, L. Ciotti, G. Turchetti, Fidelity and reversibility in the restricted three body problem. *Commun. Nonlinear Sci. Numer. Simul.* **35**, 53–68 (2016)
2. F. Panichi, K. Goździewski, G. Turchetti, The reversibility error method (REM): a new, dynamical fast indicator for planetary dynamics. *MNRAS* **468**(1), 469–491 (2017)
3. E. Lega, M. Guzzo, C. Froeschlé, Theory and applications of the fast Lyapunov indicator (FLI) Method. In: C. Skokos, G.A. Gottwald, J. Laskar (Eds) *Chaos Detection and Predictability*, pp. 35–54, Berlin, Heidelberg, 2016. Springer Berlin Heidelberg
4. M. Guzzo, E. Lega, Theory and applications of fast Lyapunov indicators to model problems of celestial mechanics. *Celestial Mech. Dyn. Astron.* **135**(4), 37 (2023)
5. M. Giovannozzi, W. Scandale, E. Todesco, Prediction of long-term stability in large hadron colliders. *Part. Accel.* **56**, 195 (1997)
6. K. Hwang, C. Mitchell, R. Ryne, Rapidly converging chaos indicator for studying dynamic aperture in a storage ring with space charge. *Phys. Rev. Accel. Beams*, **23**(8), 084601 (2020)
7. Y. Li, Y. Hao, K. Hwang, R. Rainer, A. He, A. Liu, Fast dynamic aperture optimization with forward-reversal integration. *Nucl. Instrum. Methods Phys. Res., Sect. A* **988**, 164936 (2021)
8. C.E. Montanari, A. Bazzani, M. Giovannozzi, G. Turchetti, Using dynamic indicators for probing single-particle stability in circular accelerators. In: *Proceedings of the IPAC'22*, number 13 in international particle accelerator conference, pp. 168–171. JACoW Publishing, Geneva, Switzerland, 07 2022
9. E. Todesco, M. Giovannozzi, Dynamic aperture estimates and phase-space distortions in nonlinear betatron motion. *Phys. Rev. E*, **53**(4), 4067–4076 (1996)
10. M. Giovannozzi, W. Scandale, E. Todesco, Dynamic aperture extrapolation in presence of tune modulation. *Phys. Rev.*, **E57**(3), 3432 (1998)
11. A. Bazzani, M. Giovannozzi, E. H. Maclean, C. E. Montanari, F. F. Van der Veken, W. Van Goethem, Advances on the modeling of the time evolution of dynamic aperture of hadron circular accelerators. *Phys. Rev. Accel. Beams*, **22**, 104003 (2019)
12. O.S. Brüning, P. Collier, Ph. Lebrun, S. Myers, R. Ostojic, J. Poole, P. Proudlock, LHC Design Report (Monographs. CERN, Geneva, CERN Yellow Reports, 2004)
13. O. Aberle et al. High-Luminosity Large Hadron Collider (HL-LHC): Technical design report. CERN Yellow Reports: Monographs. CERN, Geneva (2020)
14. G. Arduini, J. Barranco, A. Bertarelli, N. Biancacci, R. Bruce, O. Brüning, X. Buffat, Y. Cai, L.R. Carver, S. Fartoukh, M. Giovannozzi, G. Iadarola, K. Li, A. Lechner, L. Medina Medrano, E. Métral, Y. Nosochkov, Y. Papaphilippou, D. Pellegrini, T. Pieloni, J. Qiang, S. Redaelli, A. Romano, L. Rossi, G. Rumolo, B. Salvant, M. Schenk, C. Tambasco, R. Tomás, S. Valishev, F.F. Van der Veken, High luminosity LHC: challenges and plans. *J. Instrument.* **11**(12), C12081 (2016)
15. S. Das, Y. Saiki, E. Sander, J.A. Yorke, Quantitative quasiperiodicity. *Nonlinearity*, **30**(11):4111–4140 (2017)
16. S. Das, J.A. Yorke, Super convergence of ergodic averages for quasiperiodic orbits. *Nonlinearity*, **31**(2), 491–501 (2018)
17. F. Panichi, L. Ciotti, G. Turchetti, Fidelity and reversibility in the restricted three body problem. *Commun. Nonlinear Sci. Numer. Simul.* **35**, 53–68 (2016)
18. J. Laskar, Introduction to frequency map analysis. In: C. Simó (Ed) *Hamiltonian Systems with Three or More Degrees of Freedom*, pp. 134 — 150, Dordrecht. Springer, Springer (1999)
19. J. Laskar, Frequency map analysis and quasiperiodic decompositions. [arXiv: Dynamical Systems](https://arxiv.org/abs/2003.01401) (2003)
20. J. Laskar, Frequency map analysis of an Hamiltonian system. In: *AIP conference proceedings*, vol. 344, pp. 130–159. American Institute of Physics (1995)
21. E. Lega, C. Froeschlé, Numerical investigations of the structure around an invariant KAM torus using the frequency map analysis. *Phys. D* **95**(2), 97–106 (1996)
22. Y. Papaphilippou, J. Laskar, Frequency map analysis and global dynamics in a galactic potential with two degrees of freedom. *Astron. Astrophys.* **307**, 427–449 (1996)
23. Y. Papaphilippou, J. Laskar, Global dynamics of triaxial galactic models through frequency map analysis. *Astron. Astrophys.* **329**, 451–481 (1998)
24. Y. Papaphilippou, Global dynamics of a galactic potential via frequency map analysis. In C. Simó (Ed), *Hamiltonian Systems with Three or More Degrees of Freedom*, pp. 523–527, Dordrecht. Springer Netherlands (1999)
25. J. Laskar, Application of frequency map analysis. In: *The Chaotic Universe: Proceedings of the Second ICRA Network Workshop*, Rome, Pescara, Italy, 1–5 February 1999, vol. 10, pp. 115. World Scientific (2000)
26. M. Comunian, A. Pisent, A. Bazzani, G. Turchetti, S. Rambaldi, Frequency map analysis of a three-dimensional particle in the core model of a high intensity linac. *Phys. Rev. ST Accel. Beams* **4**, 124201 (2001)
27. J. Laskar, Frequency map analysis and particle accelerators. In: *Proceedings of the 2003 Particle Accelerator Conference*, vol. 1, pp. 378–382 (2003)
28. Y. Papaphilippou, L. Farvacque, J. Laskar, A. Ropert, Probing the non-linear dynamics of the ESRF storage ring with experimental frequency maps. In: *Proceedings of the 2003 Particle Accelerator Conference*, vol. 1, pp. 3189–3191. IEEE (2003)
29. L. Nadolski, J. Laskar, Review of single particle dynamics for third generation light sources through frequency map analysis. *Phys. Rev. ST Accel. Beams*, **6**, 114801 (2003)
30. T. Shun-Qiang, L. Gui-Min, L. Hao-Hu, H. Jie, C. Guang-Ling, W. Cheng-Lan, Nonlinear optimization of the modern synchrotron radiation storage ring based on frequency map analysis. *Chin. Phys. C* **33**(2), 127 (2009)
31. D. Shatilov, E. Levichev, E. Simonov, M. Zobov, Application of frequency map analysis to beam-beam effects study in crab waist collision scheme. *Phys. Rev. ST Accel. Beams*, **14**, 014001 (2011)
32. Y. Papaphilippou, Detecting chaos in particle accelerators through the frequency map analysis method. *Chaos* **24**(2), 024412 (2014)
33. T. Tydecks et al. FCC-ee Dynamic aperture studies and frequency map analysis. In *Proceedings of the IPAC'18*, pp. 244–246. JACoW Publishing, Geneva, Switzerland (2018)
34. P. Zisopoulos, Y. Papaphilippou, J. Laskar, Refined betatron tune measurements by mixing beam position data. *Phys. Rev. Accel. Beams*, **22**, 071002 (2019)
35. A. Bazzani, M. Giovannozzi, C. E. Montanari, G. Turchetti, Performance analysis of indicators of chaos for nonlinear dynamical systems. *Phys. Rev. E*, **107**, 064209 (2023)

36. A. Bazzani, G. Servizi, E. Todesco, G. Turchetti, A Normal Form Approach to the Theory of Nonlinear Betatronic Motion (Monographs. CERN, Geneva, CERN Yellow Reports, 1994)
37. A. Bazzani, F. Brini, G. Turchetti, Diffusion of the adiabatic invariant for modulated symplectic maps. *AIP Conf. Proc.* **395**, 129 (1997)
38. A. Bazzani, S. Siboni, G. Turchetti, Diffusion in Hamiltonian systems with a small stochastic perturbation. *Phys. D* **76**(1), 8–21 (1994)
39. C. Froeschlé, E. Lega. Weak chaos and diffusion in Hamiltonian systems. In: *The Dynamics of Small Bodies in the Solar System*, pp. 463–502. Springer, Berlin, Heidelberg (1999)
40. A. Morbidelli, C. Froeschlé, On the relationship between lyapunov times and macroscopic instability times. *Celestial Mech. Dyn. Astron.*, **63**(2), 227–239 (1995)
41. A. Morbidelli, A. Giorgilli, On a connection between KAM and Nekhoroshev's theorems. *Phys. D* **86**(3), 514–516 (1995)
42. P.M. Cincotta, C.M. Giordano, I.I. Shevchenko, Revisiting the relation between the lyapunov time and the instability time. *Phys. D* **430**, 133101 (2022)
43. M. Giovannozzi, E. McIntosh, Development of parallel codes for the study of nonlinear beam dynamics. *Int. J. Mod. Phys. C*, **8**(11) 155–170 (1996)
44. J. Ghorpade, J. Parande, M. Kulkarni, A. Bawaskar, GPGPU Processing in CUDA Architecture. *CoRR*, abs/1202.4347 (2012)
45. G. Iadarola, R. De Maria, S. Lopaciuk, A. Abramov, X. Buffat, D. Demetriadou, L. Deniau, P. Hermes, P. Kicsiny, P. Kruyt, A. Latina, L. Mether, K. Paraschou, Sterbini, F. Van Der Veken, P. Belanger, P. Niedermayer, D. Di Croce, T. Pieloni, L. Van Riesen-Haupt. Xsuite: an integrated beam physics simulation framework (2023)
46. G. Iadarola. Xsuite documentation (2022). <https://xsuite.readthedocs.io>
47. R. De Maria, J. Andersson, V.K. Berglyd Olsen, L. Field, M. Giovannozzi, P.D. Hermes, N. Høimyr, G. Iadarola, S. Kostoglou, E. H. Maclean, E. McIntosh, A. Mereghetti, J. Molson, D. Pellegrini, T. Persson, M. Schwinzerl, B. Dalena, T. Pugnat, I. Zacharov, K.N. Sjobak. SixTrack version 5: status and new developments. *J. Phys. Conf. Ser.* **1350**(11), 012129 (2019)
48. X. Pang, L. Rybarczyk. GPU accelerated online multi-particle beam dynamics simulator for ion linear particle accelerators. *Comput. Phys. Commun.* **185**(3), 744–753 (2014)
49. A. Oeftiger, S. Hegglin. Space Charge Modules for PyHEADTAIL. In *Proceedings HB'16*, pp. 124–129. JACoW Publishing, Geneva, Switzerland (2016)
50. A. Adelman, P. Calvo, M. Frey, A. Gsell, U. Locans, C. Metzger-Kraus, N. Neveu, C. Rogers, S. Russell, S. Sheehy, et al, OPAL a versatile tool for charged particle accelerator simulations. *arXiv preprint arXiv:1905.06654* (2019)
51. M. Schwinzerl, H. Bartosik, R. De Maria, G. Iadarola, A. Oeftiger, K. Paraschou, Optimising and extending a single-particle tracking library for high parallel performance. In *Proceedings of the IPAC'21*, pp 4146–4149. JACoW Publishing, Geneva, Switzerland (2021)
52. P.D. Hermes, R. Bruce, R. De Maria, M. Giovannozzi, G. Iadarola, D. Mirarchi, S. Redaelli. A Novel Tool for Beam Dynamics Studies with Hollow Electron Lenses. In *Proceedings of the IPAC'22*, number 13 in international particle accelerator conference, pp. 176–179. JACoW Publishing, Geneva, Switzerland, 07 2022
53. K. Iliakis, H. Timko, S. Xydis, P. Tsapatsaris, D. Soudris, Enabling large scale simulations for particle accelerators. *IEEE Trans. Parallel Distrib. Syst.* **33**(12), 4425–4439 (2022)
54. L. van Riesen-Haupt et al., Benchmarking equilibrium emittance simulation tools for the future circular collider. In *Proceedings of the IPAC'24*, number 15 in IPAC'24—15th international particle accelerator conference, pp. 2461–2464. JACoW Publishing, Geneva, Switzerland, 05 2024
55. A. Fornara et al., Emittance growth studies due to crab cavity induced amplitude noise in the sps. In: *Proceedings of the IPAC'24*, number 15 in IPAC'24—15th International Particle Accelerator Conference, pp. 3163–3166. JACoW Publishing, Geneva, Switzerland, 05 2024
56. Ch. Skokos, The Lyapunov characteristic exponents and their computation. In: J.-J. Souchay, R. Dvorak (Eds) *Dynamics of Small Solar System Bodies and Exoplanets*, pp. 63–135. Berlin, Heidelberg, 2010. Springer Berlin Heidelberg
57. R. Bartolini, A. Bazzani, M. Giovannozzi, W. Scandale, E. Todesco. Tune evaluation in simulations and experiments. *Part. Accel.*, **52**, 147–177 (1995)
58. R. Bartolini, M. Giovannozzi, W. Scandale, A. Bazzani, E. Todesco, Precise measurement of the betatron tune. *Part. Accel.* **55**, 1–10 (1996)
59. G. Russo, G. Franchetti, M. Giovannozzi. New techniques to compute the linear tune. In *Proceedings IPAC'21*, number 12 in International Particle Accelerator Conference, pp. 4142–4145, Geneva, 08 2021. JACoW Publishing, Geneva, Switzerland
60. R. De Maria. HL-LHC Optics v1.5 Update, (2019)
61. R. De Maria. HL-LHC MAD-X optics V1.5 Tag v0.1, (2020)
62. MAD - Methodical Accelerator Design. <https://mad.web.cern.ch/mad/>
63. G. Tancredi, A. Sánchez, F. Roig. A comparison between methods to compute lyapunov exponents. *Astron. J.* **121**(2), 1171 (2001)
64. A.N. Kolmogorov, On the Conservation of Conditionally Periodic Motions under Small Perturbation of the Hamiltonian. *Dokl. Akad. Nauk SSR* **98**, 527 (1954)
65. J. Moser. On invariant curves of area-preserving mappings of an annulus. *Nachr. Akad. Wiss. Göttingen Math.-Phys. Kl., II* (1962)
66. V.I. Arnold. Proof of a theorem of A.N. Kolmogorov on the preservation of conditionally periodic motions under a small perturbation of the Hamiltonian. *Russ. Math. Surv.* **18**, 9 (1963)
67. N. Nekhoroshev. Behavior of Hamiltonian systems close to integrable. In: *Functional Analysis and Its Applications*, vol. 5, pp. 338. Kluwer Academic Publishers-Plenum Publishers (1971)
68. N. Nekhoroshev. An exponential estimate of the time of stability of nearly-integrable Hamiltonian systems. *Russ. Math. Surv.*, **32**(6), 1 (1977)
69. A. Bazzani, S. Marmi, G. Turchetti, Nekhoroshev estimate for isochronous non resonant symplectic maps. *Cel. Mech.* **47**(4), 333 (1990)
70. G. Turchetti. Nekhoroshev stability estimates for symplectic maps and physical applications. In J.-M. Luck, P. Moussa, M. Waldschmidt, (Eds) *Number Theory and Physics*, volume 47 of Springer Proceedings in Physics, pp. 223–234, Berlin, Heidelberg (1990). Springer Berlin Heidelberg
71. M. Giovannozzi. A proposed scaling law for intensity evolution in hadron storage rings based on dynamic aperture variation with time. *Phys. Rev. ST Accel. Beams*, **15**, 024001 (2012)
72. R.M. Corless, G.H. Gonnet, D.E.G. Hare, D.J. Jeffrey, D.E. Knuth. On the lambertw function. *Adv. Computation. Math.*, **5**(1), 329–359 (1996)
73. W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd edn. (Cambridge University Press, USA, 2007)
74. A. Bazzani, M. Giovannozzi, E.H. Maclean, Analysis of the non-linear beam dynamics at top energy for the CERN Large Hadron Collider by means of a diffusion model. *Eur. Phys. J. Plus* **135**(1), 77 (2020)
75. C.E. Montanari, A. Bazzani, M. Giovannozzi, Probing the diffusive behaviour of beam-halo dynamics in circular accelerators. *Eur. Phys. J. Plus* **137**(11), 1264 (2022)
76. C.E. Montanari, Diffusive approach for non-linear beam dynamics in a circular accelerator. Master thesis, Department of Physics and Astronomy, Alma Mater Studiorum—University of Bologna, 2019. CERN-THESIS-2019-383
77. C.E. Montanari, A. Bazzani, M. Giovannozzi, A.A. Gorzawski, S. Redaelli, Testing the Global Diffusive Behaviour of Beam-Halo Dynamics at the CERN LHC Using Collimator Scans. In *Proc. IPAC'22*, number 13 in International Particle Accelerator Conference, pp. 172–175. JACoW Publishing, Geneva, Switzerland, (2022)

78. C.E. Montanari, R.B. Appleby, A. Bazzani, M. Giovannozzi, P. Hermes, A. Poyet, S. Redaelli, G. Sterbini, *Measurement of the Nonlinear Diffusion of the Proton Beam Halo at the CERN LHC* (Accepted for publication in Eur. Phys. J. Plus, 2025)
79. A. Milani, A.M. Nobili, An example of stable chaos in the solar system. *Nature* **357**(6379), 569–571 (1992)
80. A. Milani, A.M. Nobili, Z. Knežević, Stable chaos in the asteroid belt. *Icarus* **125**(1), 13–31 (1997)
81. S. Elaskar, E. del Río. Review of chaotic intermittency. *Symmetry*, **15**(6) (2023)
82. C.H. Skokos, G. A Gottwald, J. Laskar. *Chaos Detection and Predictability*, vol. 1. Springer (2016)
83. C. Froeschlé, R. Gonczi, E. Lega. The fast Lyapunov indicator: a simple tool to detect weak chaos. Application to the structure of the main asteroidal belt. *Planet. Space Sci.* **45**(7), 881–886 (1997). Asteroids, Comets, Meteors 1996 - II
84. R. Barrio, Theory and applications of the orthogonal fast lyapunov indicator (OFLI and OFLI2) methods. *Chaos Detect. Predictab.* **915**, 55–92 (2016)
85. F. Panichi, G. Turchetti, Lyapunov and reversibility errors for Hamiltonian flows. *Chaos Solitons Fractals* **112**, 83–91 (2018)
86. G. Turchetti, F. Panichi, Fast indicators for orbital stability: a survey on lyapunov and reversibility errors, in *Progress in Relativity, chapter 10*. ed. by C.G. Buzea, M. Agop, L. Butler (IntechOpen, Rijeka, 2019)
87. IEEE Standard for Floating-Point Arithmetic. IEEE Std 754-2019 (Revision of IEEE 754-2008), pp. 1–84, (2019)
88. J. Pöschel, The concept of integrability on cantor sets for Hamiltonian systems. *Celestial Mech. Dyn. Astron.* **28**(1), 133–139 (1982)
89. J. Laskar, C. Froeschlé, A. Celletti. The measure of chaos by the numerical analysis of the fundamental frequencies. *Appl. Stand. Map. Phys. D*, **56**(2–3), 253–269 (1992)
90. R. Bartolini, F. Schmidt. A computer code for frequency analysis of non-linear betatron motion. Technical report, CERN-SL-Note-98-017-AP (1998)