

Geometry of almost-conserved quantities in symplectic maps: Approximate invariants in nonlinear accelerator systems

T. Zolkin[†]*Independent Researcher, Atlanta, Georgia 30341, USA*S. Nagaitsev[‡]*Brookhaven National Laboratory, Upton, New York 11973, USA
and Old Dominion University, Norfolk, Virginia 23529, USA*I. Morozov^{*,§}*Novosibirsk State Technical University, Novosibirsk 630073, Russia*S. Kladov^{||}*University of Chicago, Chicago, Illinois 60637, USA*

(Received 12 May 2025; accepted 27 April 2026; published 29 June 2026)

We present a perturbative method for constructing approximate invariants of motion directly from the equations of discrete-time symplectic systems. This framework offers a natural nonlinear extension of the classic Courant-Snyder (CS) theory for systems with 1 degree of freedom—a foundational cornerstone in accelerator physics now spanning seven decades and historically focused on linear phenomena. The original CS formalism emerged under conditions where nonlinearities were weak, design goals favored linear motion, and analytical tools—such as the Kolmogorov-Arnold-Moser theory—had not yet been fully developed. While various normal-form methods have been proposed to treat near-integrable dynamics, the approach introduced here stands out for its conceptual transparency, minimal computational overhead, and direct applicability to realistic systems. We demonstrate its power and versatility by applying it to several operational accelerator configurations at the Fermi National Accelerator Laboratory (Fermilab), illustrating how the method enables fast, interpretable diagnostics of nonlinear behavior across a broad range of machine conditions.

DOI: 10.1103/zxgv-2xf

Introduction. Since its inception, the foundational theory of motion in circular particle accelerators has been built around linear dynamics. The celebrated Courant-Snyder (CS) theory [1], developed in the 1950s, provided a powerful analytical framework for understanding phase space evolution in systems weakly perturbed from harmonic oscillations, forming the theoretical backbone of modern accelerator design [2,3]. However, this linear formalism is increasingly strained in the face of modern challenges: higher beam intensities, smaller beam emittance, tighter stability margins, and the pervasive presence

of nonlinear magnetic elements in advanced accelerator facilities such as the Integrable Optics Test Accelerator (IOTA) [4,5], as well as in proposed next-generation machines like the muon collider [6] or the FCC [7–9].

Despite the critical role of nonlinear effects in determining long-term stability and beam quality, a general analytical extension of the Courant-Snyder formalism to strongly nonlinear regimes has remained elusive, even for 1 degree of freedom. In particular, the twist expansion of action-angle variables achieved via Lie-algebraic methods is compatible with numerous invariants of motion, all of which are conserved to the same order of approximation (see Ref. [10] for details). In the widely used formalism of Birkhoff normal forms, this ambiguity is commonly addressed by isolating resonances, which requires the manual selection of a specific normal form. In this work, we present an approach that extends this understanding: a systematic perturbative framework for constructing approximate invariants in symplectic, reversible maps [11], in which a different averaging strategy allows multiple resonances to be retained up to a given perturbative order. We note that: (i) direct comparisons of our method with

*Also at Elettra Sincrotrone Trieste, Trieste, Italy.

[†]Contact author: iguanodyn@gmail.com

[‡]Contact author: snagaitsev@bnl.gov

[§]Contact author: i.morozov@corp.nstu.ru

^{||}Contact author: kladov@uchicago.edu

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

resonant normal forms (isolated-resonance approximation) can be made by juxtaposing Fig. 20 in [10] with Fig. 6 of [12]; and (ii) our averaging technique recovers exact invariants in integrable systems, providing a natural benchmark for its applicability and demonstrating consistency with established results. Prominently, the method is compact both in formulation and implementation, relying on a single guiding principle: approximate invariance.

Consider a symplectic transformation $T: \mathbb{R}^2 \mapsto \mathbb{R}^2$ describing a dynamical system with 1 degree of freedom, characterized by canonical coordinate q and conjugate momentum p . In the integrable case, an exact invariant $\mathcal{K}[p, q]$ exists and satisfies

$$\mathcal{K}[T(p, q)] - \mathcal{K}[p, q] = 0, \quad \forall (q, p) \in \mathbb{R}^2.$$

We generalize this to a nonintegrable scenario by relaxing the right-hand side to allow deviations of order $\mathcal{O}(\epsilon^{n+1})$, where n defines the order of perturbation. From this single condition, we construct a parameterized function that captures the essential geometry of phase space—revealing the topology of nonlinear motion, the structure of island chains, separatrices, and the amplitude dependence of the rotation number. As discussed in the first part of this series [10], the resulting approximate invariant also preserves two families of reversibility-induced symmetries up to the same order, reinforcing both its geometric fidelity and consistency.

The resulting method is both powerful and remarkably efficient, suitable for both numerical and analytical computations. All examples presented in this work—spanning multiple accelerators from the Fermilab facility—are generated from a single, closed-form function parameterized by the elements of the accelerator lattice. More precisely, it depends only on the coefficients of the discrete-time transformation (the *one-turn map*) that describes the evolution of a particle over a single revolution with respect to a reference trajectory. This function acts as a nonlinear generalization of the Courant-Snyder invariant, enabling rapid visualization of stability boundaries, identification of resonant structures, and quantitative prediction of long-term dynamics—without requiring particle tracking or extensive numerical simulation. When combined with the Danilov’s theorem [13,14], this technique can be connected to standard methods, allowing the construction of a set of approximate action-angle variables and nonlinear betatron tune.

In the first part of this series [10], we developed the theoretical foundations of our perturbative approach by introducing the central concept of an approximate invariant constructed order by order in the small parameter ϵ . This parameter naturally emerges from the smallness of oscillation amplitudes and does not, in itself, require the nonlinear forces to be perturbatively weak. We demonstrated that these invariants respect the discrete reversibility symmetries—serving as a discrete analog of Noether’s theorem—and benchmarked the results against standard

Lie-algebraic normal-form methods near a fixed point [15–17]. We also showed that, at each perturbative order, the approximate invariant is not unique but is determined only up to arbitrary terms composed of powers of the zeroth-order (linear) Courant-Snyder invariant (see Ref. [10] for details). Regardless of the choice of these coefficients, the twist expansion obtained from the standard Lie normal-form procedure remains satisfied. To resolve this arbitrariness, we introduced an averaging procedure that minimizes the fluctuations of the approximate invariant. This not only fixes the previously undetermined coefficients but also effectively resolves degeneracies associated with rational rotation numbers, making the theory applicable to singular (with diverging twist coefficients) and to n -island chain resonances. We note that this averaging differs from the traditional approach used for isolated resonances [12] and can properly account for all resonances up to a given order of approximation, thereby reinforcing both the geometric and physical foundation of the method.

The second part of [10] extends the discussion from local, infinitesimal dynamics to large-amplitude motion, aiming for a global reconstruction of (quasi-) integrable structures in both integrable and chaotic systems. Applying the method to several representative maps—including the McMillan, Suris exponential, and Hénon quadratic and cubic area-preserving transformations—we demonstrated both convergence and resonance compatibility. The averaged invariant not only exhibited the expected asymptotic behavior near the fixed point but also reproduced exact global results for integrable systems, successfully capturing separatrices and other invariant structures beyond the local regime (see Ref. [10] for details). A detailed comparison with the square matrix method [18,19], which is equivalent to the full normalization in the language of normal forms, showed comparable local accuracy but superior global stability and resonance performance in the present approach.

In this final part, we briefly restate the method for general planar maps and then move directly to demonstrating its practical power. We apply the framework to realistic dynamical models drawn from accelerator lattice configurations across the Fermilab complex. Specifically, we examine three major facilities: the Integrable Optics Test Accelerator (IOTA), the Main Injector, and the Mu2e Delivery Ring. These systems exhibit a wide range of nonlinear behavior, including high-order resonances and large-amplitude instabilities. In each case, our method yields rapid and interpretable diagnostics that align closely with results from full-scale particle tracking—while dramatically reducing computational demands.

Beyond its immediate practical utility, this approach offers a scalable framework for nonlinear design and optimization—potentially transforming how we diagnose, predict, and ultimately control stability in high-performance machines. It introduces a new paradigm for describing nonlinearity. Just as the linear CS theory allows

one to bypass the detailed specification of all quadrupole magnets—often numbering in the hundreds—by reducing the system to a small set of Twiss parameters that describe the one-turn map and connect directly to key dynamical quantities such as betatron tune and beam envelope, our nonlinear extension leverages the cumulative effect of all nonlinear elements, condensed into the coefficients of the one-turn map, to construct approximate invariants, thereby enabling a compact and highly informative representation of the system's dynamics.

Perturbation method. Consider a map of the plane $T: \mathbb{R}^2 \mapsto \mathbb{R}^2$, which is symplectic (i.e., it preserves area and orientation) and admits a power series expansion:

$$\begin{aligned} q' &= A_{1,0} q + A_{0,1} p + A_{2,0} q^2 + A_{1,1} q p + A_{0,2} p^2 + \cdots, \\ p' &= B_{1,0} q + B_{0,1} p + B_{2,0} q^2 + B_{1,1} q p + B_{0,2} p^2 + \cdots. \end{aligned}$$

We begin by introducing a general polynomial ansatz for the approximate invariant of order n :

$$\mathcal{K}^{(n)}[p, q] = \mathcal{K}_0 + \epsilon \mathcal{K}_1 + \cdots + \epsilon^n \mathcal{K}_n,$$

where $\mathcal{K}_m$ is a homogeneous polynomial of degree $(m+2)$:

$$\mathcal{K}_m[p, q] = \sum_{\substack{i,j \geq 0 \\ i+j=m+2}} C_{i,j} p^i q^j.$$

To determine the coefficients $C_{i,j}$, we require that the approximate invariant be conserved up to $\mathcal{O}(\epsilon^{n+1})$:

$$\mathcal{R}_n = \mathcal{K}^{(n)}[p', q'] - \mathcal{K}^{(n)}[p, q] = \bar{\mathcal{R}}_n \epsilon^{n+1} + \mathcal{O}(\epsilon^{n+2}).$$

The lowest-order contribution, which can also be derived via standard linearization, is defined only up to a constant multiplier

$$\mathcal{K}_0[p, q] = C_0 [A_{0,1} p^2 + (A_{1,0} - B_{0,1}) p q - B_{1,0} q^2]. \quad (1)$$

The *seed* coefficient C_0 can be set to unity or chosen specifically to eliminate resonant denominators.

Due to underdeterminacy, the invariant (after all two-indexed coefficients $C_{i,j}$ have been fixed) is known only up to a series:

$$C_0 \mathcal{K}_0 + C_1 \mathcal{K}_1^2 + \cdots.$$

Nevertheless, the approximate invariance condition holds to the required order, and the twist coefficient τ_k converges at order $n = 2k + 2$.

The remaining coefficients C_k (for $k > 0$) are obtained via the *averaging procedure*, which minimizes the residual error. After transforming to the eigenbasis of the Jacobian, $(q, p) \rightarrow (Q, P)$, the linear part of the map becomes a pure rotation, and the zeroth-order invariant takes the form

$$\mathcal{K}_0[P, Q] = P^2 + Q^2.$$

The coefficients C_k are then determined by solving another system of linear equations

$$\frac{d}{dC_k} I_n = 0,$$

where the integral

$$I_n = \int_0^{2\pi} \overline{\mathcal{R}_n^2}[\rho, \psi] d\psi.$$

Here, (ρ, ψ) are polar phase space coordinates defined by $Q = \rho \cos \psi$ and $P = \rho \sin \psi$.

Nonlinear Courant-Snyder theory: In accelerator physics, the linear theory of motion uses a parameterization in which the linearized one-turn map (the discrete transformation that evolves a particle's phase space coordinates over one full revolution)

$$\begin{aligned} q' &= A_{1,0} q + A_{0,1} p, \\ p' &= B_{1,0} q + B_{0,1} p \end{aligned}$$

is written as

$$\begin{bmatrix} q' \\ p' \end{bmatrix} = \begin{bmatrix} \cos \Phi + \alpha \sin \Phi & \beta \sin \Phi \\ -\gamma \sin \Phi & \cos \Phi - \alpha \sin \Phi \end{bmatrix} \begin{bmatrix} q \\ p \end{bmatrix}.$$

The coefficients α , β , and γ , collectively known as the *Courant-Snyder/Twiss parameters*, are not all independent: the symplectic condition imposes the constraint $\beta\gamma = 1 + \alpha^2$; conventionally, the *beta function* is chosen to be positive, $\beta > 0$. $\Phi = 2\pi\nu_0$ is the *betatron phase advance*, with ν_0 being the bare *betatron tune*—the unperturbed rotation number.

This parameterization enables fast, linearized estimates of several experimentally important quantities. First, (i) one obtains the betatron tune, corresponding to the frequency of horizontal or vertical oscillations around the closed reference orbit (or the synchrotron tune in the longitudinal case):

$$\nu_0 = \frac{1}{2\pi} \arccos \frac{A_{1,0} + B_{0,1}}{2}.$$

Second, (ii) one obtains the *Courant-Snyder invariant*

$$\mathcal{K}^{(0)}[p, q] = \gamma q^2 + 2\alpha q p + \beta p^2,$$

which provides a convenient parameterization of trajectories in the phase space coordinates at a fixed location in the accelerator. By propagating the Twiss parameters along the reference orbit, this invariant allows the reconstruction of phase space trajectories not only at each discrete turn but also continuously along the machine's circumference, s .

This invariant can also be obtained using Eq. (1) by selecting $C_0 = \beta^*/A_{0,1}$, where $\beta^* = \beta(s^*)$ is the beta

function at a reference location, s^* . Moreover, (iii) this scaling of the invariant,

$$\mathcal{K}^{(0)}[p, q] = C_0 \mathcal{K}_0 = \text{em}$$

provides a direct correspondence between the invariant and the *beam emittance* em , which quantifies the area occupied by the beam in phase space. This value plays a crucial role in estimating performance metrics of colliders such as luminosity and in machines involving strong dissipative forces, including a potential muon collider [6] and facilities employing electron or optical stochastic cooling [5]. The perturbative method described earlier naturally extends this linear framework.

The zeroth-order Courant-Snyder invariant is generalized to a sequence of approximate invariants

$$\mathcal{K}^{(0)} \rightarrow \mathcal{K}^{(1)} \rightarrow \mathcal{K}^{(2)} \rightarrow \dots$$

In many practical cases, low-order approximations suffice to accurately describe particle motion, as seen in Fig. 1.

Similarly, the linear tune ν_0 is replaced by its nonlinear analogs:

$$\nu_0 \rightarrow \nu_1 \rightarrow \nu_2 \rightarrow \dots$$

These are extracted from the approximate invariants using Danilov's theorem [13,14]

$$\nu_n = \int_q^{q'} \left(\frac{\partial \mathcal{K}^{(n)}}{\partial p} \right)^{-1} dq / \oint \left(\frac{\partial \mathcal{K}^{(n)}}{\partial p} \right)^{-1} dq$$

or equivalently,

$$\nu_n = dJ'_n / dJ_n,$$

where J_n and J'_n are *action* and *partial action* variables:

$$J_n = \oint \frac{p dq}{2\pi}, \quad \text{and} \quad J'_n = \int_q^{q'} \frac{p dq}{2\pi}.$$

All integrals are taken along a constant level set $\mathcal{K}^{(n)}[p, q] = \text{const}$. The integral $\int_q^{q'}$ spans one iteration of the map (independent of the choice of q), while $\oint$ denotes a full contour around a closed trajectory.

This series of approximations differs from the traditional *twist expansion*:

$$\nu(J) = \nu_0 + \tau_0 J + \frac{1}{2!} \tau_1 J^2 + \frac{1}{3!} \tau_2 J^3 + \dots$$

At the same time, a comparison between the twist expansion and the n th order approximation series

$$\nu_n(J) = \nu_0 + \tau_0^{(n)} J + \frac{1}{2!} \tau_1^{(n)} J^2 + \frac{1}{3!} \tau_2^{(n)} J^3 + \dots \quad (2)$$

shows that for each even order $n = 2m$, all coefficients up to $\tau_{m-1}^{(n)}$ converge to their exact values τ_{m-1} .

As n increases, ν_n approaches the twist expansion $\nu(J)$, Eq. (2), and both approximations converge to the integrable limit of the Twist map. Yet, for nonintegrable systems,




FIG. 1. Comparison between numerically computed trajectories (black dots) and constant level sets of the approximate invariant (red curves) for various accelerator lattices at Fermilab. Panels (a1) and (a2) correspond to different configurations of slow extraction in the Delivery Ring for the Mu2e experiment, panel (b) corresponds to the IOTA ring, and panels (c1) and (c2) correspond to the Main Injector ring. All plots display the horizontal phase space, with position q in millimeters and conjugate momentum $p = dq/ds$; for panel (b), both q and p are normalized to the beam's rms size (in units of beam σ).

this convergence comes at the cost of a vanishingly small domain. At the other extreme, when n is small but nonzero, the two methods differ in form but provide complementary insights into the system's behavior. The low-order *twist coefficients* like

$$\tau_0 = \frac{d\nu}{dJ}, \quad \text{and} \quad \tau_1 = \frac{d^2\nu}{dJ^2},$$

capture key nonlinear features and are widely used in estimates of collective beam stability [20,21]. At the same time, as demonstrated in the first part of this manuscript [10] and in related studies [22,23], large-amplitude behavior and the edge of stability can be significantly better captured using low-order nonlinear estimates ν_1 or ν_2 (labeled as SX-1 and SX-2, respectively), particularly near low-order singular (i.e., $\tau_0 = \infty$) resonances such as integer, half-, third-, or quarter-integer.

In the second part of [10], we also show that this method applies to the simply connected regions bounded by resonance island chains appearing in nonsingular bifurcations.

Moreover, the approximate invariant can always be rescaled to match the action, $\mathcal{K}^{(n)}[p, q] = J_n$, thus yielding a nonlinear analog of emittance, just as in the linear case.

Before proceeding with examples, we note that the transformation to the eigenbasis $(q, p) \rightarrow (Q, P)$ in Courant-Snyder theory is achieved via the *Floquet transformation* and corresponding *betatron amplitude matrix*

$$\begin{bmatrix} q \\ p \end{bmatrix} = \mathbf{B} \begin{bmatrix} Q \\ P \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \sqrt{\beta} & 0 \\ -\alpha/\sqrt{\beta} & 1/\sqrt{\beta} \end{bmatrix}.$$

Once the coefficients C_k are determined, they can be substituted into the nonaveraged invariant $\mathcal{K}^{(n)}[p, q]$, expressed directly in physical coordinates. This makes the formalism particularly convenient for experimental applications. The approximate invariant with substituted coefficients C_k is referred to as the *averaged invariant* and denoted by $\langle \mathcal{K}^{(n)} \rangle$.

Figure 1 illustrates this method for several circular accelerators at Fermilab, which differ significantly in both design and purpose: third-integer resonant slow extraction in the Delivery Ring machine for the Mu2e experiment [panels (a1) and (a2)], the Integrable Test Optics Accelerator (IOTA) ring (b), and the Main Injector ring [(c1) and (c2)]. In each case, we use the full machine map—including all nonlinear (thin or thick) elements—derived from a lattice file describing magnet locations and strengths, just as in standard accelerator simulation codes.

We compare direct numerical application of the map (black dots) with the level sets of the averaged approximate invariant $\langle \mathcal{K}^{(n)} \rangle[p, q]$ (red curves). Each red curve corresponds to the trajectory seeded by the matching black point. If a red curve is not visible, it is fully overlaid by the trajectory, emphasizing the excellent agreement.

In examples (a1) and (a2), dynamics are dominated by a singular third-order resonance, and even a low-order

approximation ($n = 4$, corresponding to J^2) gives outstanding agreement. In the IOTA ring, we observe a nonsingular third-order resonance dominated by octupole-type nonlinearity, characterized by a distinctive chain of islands. In (c), the Main Injector ring exhibits high-order resonances and chaotic behavior, but even a sixth-order approximation ($n = 6$, or J^3) accurately captures the stability region and general structure of the phase space.

Finally, we stress that the examples shown rely on unrefined, bare lattice configurations, which include imperfections and design features not yet resolved or optimized experimentally. These cases often prompt direct simulations—aimed at improving performance [as in (a1) and (a2)] or reducing detrimental resonance structures [as in (c)]. That our method performs well in such cases underscores its value as a robust and complementary tool for nonlinear dynamics analysis.

Conclusions and possible generalizations. While Fig. 1 demonstrates the power and versatility of our method through applications to three distinct accelerator rings at the Fermilab, it is important to comment on its limitations and potential extensions. First and foremost, this method is not intended to replace traditional simulations during the final stages of machine design, where collective effects must be accounted for. Nevertheless, even in such regimes, it offers critical insight into intrinsic single-particle dynamics—such as amplitude-dependent tune shifts—which are essential for estimating instability thresholds and determining the level of Landau damping required for instability suppression [21].

A natural extension of this framework is to systems with more than 1 degree of freedom, i.e., higher-dimensional symplectic maps. While a treatment of nonlinear coupling is currently under consideration, we note that even the theory of coupled linear motion was historically developed significantly after the original Courant-Snyder framework (see, e.g., coupled betatron theory [24] for 2 degrees of freedom and [25] for 3 degrees). Therefore, a few brief comments are in order. In principle, the same perturbative approach can be applied; however, for a system with 2 degrees of freedom, one must construct two approximate invariants, each depending on four variables in a four-dimensional phase space. The trajectory then lies on a torus defined by the intersection of two three-dimensional hypersurfaces, significantly complicating both visualization and extraction of dynamical features. Still, recent extensions of the Danilov theorem to higher dimensions [26] offer a promising framework to be used in tandem with the construction of such invariants.

Despite these challenges, we have shown that even a two-dimensional analysis often yields remarkable value. This includes a range of practical scenarios: (i) accelerators with flat beams, where the horizontal beam size is much larger than the vertical; (ii) systems dominated by a strong one-dimensional resonance, as in the case of the Main

Injector; (iii) machines designed to exploit nonlinear instability for slow resonant extraction, such as the Mu2e Delivery Ring; and (iv) round beams, where this method can be directly applied to radial degree of freedom—see Ref. [27] for further discussion.

The developed procedure is conceptually simple and grounded solely in the first principle of approximate invariance. In contrast to Lie-algebraic normal-form methods, it does not rely on the exponential operator formalism and assumptions on the final normalized system. Instead, it provides a direct and geometrically transparent construction of the invariant itself, order by order, making the method physically intuitive and highly regular in application.

The framework is intrinsically linked to reversible symmetries: the invariant is sought as a function that remains nearly unchanged under two reversing symmetries up to a prescribed perturbative order. In this sense, the method parallels a discrete Noether-like approach, where reversibility substitutes for continuous symmetries.

While the commonly used full normalization procedure [16,28–30] primarily targets the computation of the amplitude-dependent tune shifts and resonance driving terms, the presented formulation succeeds in both computing the tune shifts by employing the Danilov theorem [13,14] and describing the stability boundary itself [10]. Since the full normalization tries to construct a conjugation to an integrable system depending solely on the action variables, its application to the problem of describing the stability boundary is not well grounded. Any resonant structure, like a chain of islands, inside the stability domain constitutes a boundary that cannot be crossed and leads to divergence of the complete normalization procedure. The latter also implies the presented theory can produce a more generic tune dependence picture that can cross the inner resonance chains. In attempting to describe the stability boundary, the isolated resonance model can be employed [12]. This requires careful selection of a resonance being studied and a procedure to switch between resonances following parameter changes. The map-based perturbation theory is free of such limitations and does not require an *a priori* selection of some resonance producing a coherent expression for the invariant valid across the full range of the intrinsic parameters.

In conclusion, by distilling the cumulative effect of all nonlinear lattice elements into a single approximate invariant, this approach offers a scalable and efficient tool for nonlinear analysis. It provides a compact, interpretable representation of complex dynamics and opens a new pathway for nonlinear optimization and control in next-generation accelerator systems, serving as both a conceptual and computational bridge between practical machine design and modern dynamical systems theory.

Acknowledgments. The authors would like to thank Taylor Nchako (Northwestern University) for carefully reading

this manuscript and for her helpful comments. The authors gratefully acknowledge Eric Stern, Robert Ainsworth, and Vladimir Nagaslaev (all from Fermilab) for providing the lattice files for the IOTA, Main Injector, and Delivery Rings, respectively. S.N. work is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Contract No. DE-AC05-06OR23177. S. K. is grateful to his supervisor, Professor Young-Kee Kim (University of Chicago), for her valuable mentorship and continuous support.

Data availability. The data that support the findings of this study are available from the corresponding author upon reasonable request.

-
- [1] Ernest D. Courant and Hartland S. Snyder, Theory of the alternating-gradient synchrotron, *Ann. Phys. (Amsterdam)* **3**, 1 (1958).
 - [2] Shyh-Yuan Lee, *Accelerator Physics*, 4th ed. (World Scientific Publishing Company, Singapore, 2018), 10.1142/11111.
 - [3] Helmut Wiedemann, *Particle Accelerator Physics*, 4th ed. (Springer, Cham, 2015), pp. XXIX, 1021, 10.1007/978-3-319-18317-6.
 - [4] V. Danilov and S. Nagaitsev, Nonlinear accelerator lattices with one and two analytic invariants, *Phys. Rev. ST Accel. Beams* **13**, 084002 (2010).
 - [5] V. Lebedev, Jonathan Jarvis, A. Romanov, D. Broemmelsiek, K. Carlson, S. Chattopadhyay, A. Dick, D. Edstrom, I. Lobach, S. Nagaitsev *et al.*, Experimental demonstration of optical stochastic cooling, *Nature (London)* **608**, 287 (2022).
 - [6] Diktys Stratakis, Essay: A path for the construction of a muon collider, *Phys. Rev. Lett.* **134**, 160001 (2025).
 - [7] Wolfgang Bartmann, Jean-Paul Burnet, Christian Carli, Antoine Chance, Paolo Craievich, Massimo Giovannozzi, Christophe Grojean, Johannes Gutleber, Klaus Hanke, Andre Henriques, Patrick Janot, Carlos Lourenco, Michelangelo Mangano, Thomas Otto, John Howard Poole, Srinji Rajagopalan, Tor Raubenheimer, Ezio Todesco, Timothy Paul Watson, and Guy Wilkinson, Future circular collider feasibility study report volume 1: Physics and experiments, 10.17181/CERN.9DKX.TDH9 (2025).
 - [8] Michael Benedikt *et al.*, Future circular collider feasibility study report volume 2: Accelerators, technical infrastructure and safety, 10.17181/CERN.EBAY.7W4X (2025).
 - [9] Johannes Gutleber, Luisa Ulrici, and Timothy Paul Watson, Future circular collider feasibility study report volume 3: Civil engineering, implementation and sustainability, 10.17181/CERN.I26X.V4VF (2025).
 - [10] Tim Zolkin, Sergei Nagaitsev, Ivan Morozov, and Sergei Kladov, Geometry of almost-conserved quantities in symplectic maps, *Chaos Solitons Fractals* **208**, 118059 (2026).
 - [11] J. A. G. Roberts and G. R. W. Quispel, Chaos and time-reversal symmetry. Order and chaos in reversible dynamical systems, *Phys. Rep.* **216**, 63 (1992).

- [12] A. Bazzani, M. Giovannozzi, G. Servizi, E. Todesco, and G. Turchetti, Resonant normal forms, interpolating Hamiltonians and stability analysis of area preserving maps, *Physica (Amsterdam)* **64D**, 66 (1993).
- [13] Sergei Nagaitsev and Timofey Zolkin, Betatron frequency and the Poincaré rotation number, *Phys. Rev. Accel. Beams* **23**, 054001 (2020).
- [14] Sergei Nagaitsev and Timofey Zolkin, Betatron frequency and the Poincaré rotation number, *Phys. Rev. Accel. Beams* **23**, 054001 (2020); *Phys. Rev. Accel. Beams* **29**, 029901 (E) (2026).
- [15] Michelotti Leo, *Intermediate Classical Dynamics with Applications to Beam Physics* (Wiley, New York, 1995), pp. 284–299.
- [16] Étienne Forest, Beam dynamics: A new attitude and framework, in *The Physics and Technology of Particle and Photon Beams* (Hardwood Academic/CRC Press, Amsterdam, The Netherlands, 1998), Vol. 8, pp. 138–141.
- [17] Berz Martin, Chapter 7—Repetitive systems, in *Modern Map Methods in Particle Beam Physics, Advances in Imaging and Electron Physics*, edited by Peter Hawkes (1999), Vol. 108, pp. 270–288, [10.1016/S1076-5670\(08\)70233-7](https://doi.org/10.1016/S1076-5670(08)70233-7).
- [18] Li Hua Yu, Analysis of nonlinear dynamics by square matrix method, *Phys. Rev. Accel. Beams* **20**, 034001 (2017).
- [19] Ivan Morozov, Square, <https://github.com/i-a-morozov/square> (2024).
- [20] Jacques Gareyte, Jean-Pierre Koutchouk, and F. Ruggiero, Landau damping dynamic aperture and octupole in LHC, Technical Report No. LHC-Project-Report-91, CERN, Geneva, 1997 [revised version number 1 submitted on August 21, 2003].
- [21] V. Shiltsev, Y. Alexahin, A. Burov, and A. Valishev, Landau damping of beam instabilities by electron lenses, *Phys. Rev. Lett.* **119**, 134802 (2017).
- [22] Tim Zolkin, Sergei Nagaitsev, and Ivan Morozov, Dynamics of McMillan mappings I. McMillan multipoles, *Physica (Amsterdam)* **476D**, 134620 (2025).
- [23] Tim Zolkin, Sergei Nagaitsev, Ivan Morozov, and Sergei Kladoy, Dynamics of McMillan mappings III. Symmetric map with mixed nonlinearity, *Nonlinear Dyn.* **114**, 635 (2026).
- [24] Valeri A. Lebedev and S. A. Bogacz, Betatron motion with coupling of horizontal and vertical degrees of freedom, *J. Instrum.* **5**, P10010 (2010).
- [25] Sergei Glukhov, Symmetry properties of a symplectic transport matrix and Twiss parameterization of a fully coupled motion, *Phys. Rev. Accel. Beams* **28**, 084001 (2025).
- [26] Chad E. Mitchell, Robert D. Ryne, Kilean Hwang, Sergei Nagaitsev, and Timofey Zolkin, Extracting dynamical frequencies from invariants of motion in finite-dimensional nonlinear integrable systems, *Phys. Rev. E* **103**, 062216 (2021).
- [27] Tim Zolkin, Brandon Cathey, and Sergei Nagaitsev, Dynamics of McMillan mappings II. Axially symmetric map, *Nonlinear Dyn.* **113**, 20253 (2025).
- [28] Alex J. Dragt, Liam M. Healy, Filippo Neri, Robert D. Ryne, David R. Douglas, and Etienne Forest, Marylie 3.0—A program for nonlinear analysis of accelerator and beamline lattices, *IEEE Trans. Nucl. Sci.* **32**, 2311 (1985).
- [29] Alex J. Dragt, *Lie Methods for Nonlinear Dynamics with Applications to Accelerator Physics* (University of Maryland, Center for Theoretical Physics, Department of Physics, College Park, 1997), <https://books.google.com/books?id=JNMScgAACAAJ>.
- [30] Martin Berz, *Modern Map Methods in Particle Beam Physics*, edited by Peter Hawkes and Martin Berz, *Advances in Imaging and Electron Physics* Vol. 108 (Elsevier, New York, 1999), pp. 1–318, [10.1016/S1076-5670\(08\)70227-1](https://doi.org/10.1016/S1076-5670(08)70227-1).