Frequency analysis and dynamic aperture studies in a low energy antiproton ring with realistic 3D magnetic fields

Lajos Bojtár◉[*]

CERN, CH-1211 Geneva 23, Switzerland

We present recent progress with our algorithm and its implementation called SIMPA described in a previous paper. The algorithm has a new and unique approach to long-term 4D tracking of charged particles in arbitrary static electromagnetic fields. Using the improvements described in this paper, we made frequency analysis and dynamic aperture studies in ELENA. The effect of the end fields and the perturbation introduced by the magnetic system of the electron cooler on dynamic aperture is shown. A special feature of this study is that we have not introduced any multipole errors into the model. The dynamic aperture calculated in this paper is the direct consequence of the geometry of the magnetic elements. Based on the results, we make a few suggestions to reduce the losses during the deceleration of the beam.

I. INTRODUCTION

The importance of fringe fields in small rings is well known and it has been taken into account for multipole magnets at various degrees for decades [1–4]. The approach to particle tracking we described in [5] naturally includes the end fields for all kinds of elements with the same treatment. The aim of this paper is two-fold. One is to apply our algorithm described in [5] on the extra low energy antiproton (ELENA) ring [6], the 30.4 meters circumference machine built at CERN to decelerate antiprotons. Another aim is to gain insight into the beam dynamics of ELENA by frequency analysis and dynamic aperture studies.

Long-term tracking of charged particles is a fundamental problem of accelerator physics, plasma physics and it is important in astrophysics. In a recent paper [5] we described a new algorithm allowing long-term symplectic integration of charged particle trajectories in arbitrary static magnetic and electric fields. We introduced several new ideas and pointed out the importance of two requirements for a physically valid description of the electromagnetic fields.

One important point is that the potentials describing static magnetic and electric fields must obey the Laplace equation everywhere in the domain of interest. This is a necessary condition to get physically valid fields from the potentials. Failing to satisfy this condition leads to spurious sources [7] which in turn leads to an energy drift during the tracking. All 3D grid methods suffer from this problem because there is no known interpolation satisfying the Laplace equation while being continuous between the 3D grid cells.

In accelerator physics, transfer map methods [8–10] are also applicable to handle realistic 3D electromagnetic fields. These are conceptually rather different and their implementations usually assume the existence of a closed orbit, which puts restrictions to the range of problems it can handle. FFAs, stellarators, tokamaks for example are difficult to handle with transfer maps, but our algorithm can treat those cases without complications.

Another approach to deal with the problem of the representation of 3D fields is to take the Fourier transform of the field in three variables [11–14]. One disadvantage of this method is the high number of harmonics needed when the field has detailed local features.

The second requirement for a physically valid description of the field is the continuity of the potentials. Cutting the potentials at some distance, even far from a magnetic element, leads to a systematic error which accumulates during the tracking leading to an energy drift and non-preservation of the phase space volume. This has been recognized also by others [2,4,15] and mitigations has been used to handle this effect. Our method avoids cutting the potentials by treating the entire volume of interest as a whole.

Before this study, little was known about the effect of the magnetic system of the electron cooler on the ELENA beam dynamics. The electron cooler is described in detail in [6]. The magnetic elements in the electron cooler having a significant effect of the beam dynamics are the following: a

---

[*]Lajos.Bojtar@cern.ch

drift solenoid with 10 mT nominal field, two toroid sections at the ends of the solenoid, two solenoid compensators, and two orbit correctors to compensate the kicks of the toroids. The components of the magnetic field of the electron cooler are shown in Fig. 1. The sharp negative spike on the B_z curve gets stronger as the distance from the design orbit gets bigger. This is due to the shielding around the cooler, which collects some field lines from the toroid coils. The effect of the toroid coils in the cooler was mostly unknown, not due to lack of effort, but due to the lack of appropriate tracking tools. A considerable amount of work went into the development of a tracking tool by collegues at CERN to address this issue [16]. The tool was based on a field map given on a 3D grid. It produced manifestly nonphysical behavior already in a few hundred turns.

This triggered the author to devise the algorithm [5] and implement it. The name of the software is SIMPA, an abbreviation of symplectic integration through monopole arrangements. At present it is a Java library. A lot of work remains to be done before it can be released for general use. It is intended to become a tool to produce field maps suitable for long-term, charged particle tracking in general, not only in the domain of accelerator physics. It contains code to use the generated field maps in actual trackings. Recent progress [17] made possible to implement an explicit second-order symplectic integrator for nonseparable Hamiltonians, which arise when magnetic fields are involved.

The simulation tools needed for frequency analysis and dynamic aperture studies described in this paper were written in Java using the SIMPA library. Both the simulation source code and the library source code is available from the author upon request. The SIMPA source code will be released under an open source license in the future. People interested in the code are encouraged to contact the author.

The algorithm described in [5] was applied to a simple test machine with a beam region having a constant circular
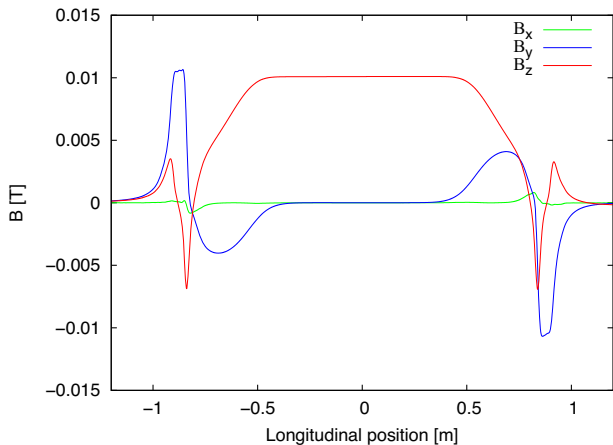

FIG. 1.   The magnetic field components of the electron cooler plotted against the longitudinal position with transverse coordinates $x = y = 2$ cm referenced to the design orbit.

profile. This was covered with a single line of overlapping spheres. See Fig. 4 in [5] for an illustration. Actual machines like ELENA have beam regions with more complex geometries. Much progress has been made to extend the algorithm to general volumes. Full description of the enhanced algorithm is outside the scope of this work and will be described separately, only a summary follows.

The paper is organized in the following way. We present the progress made to cover general volumes and explain also how this leads to gain in speed and accuracy. Then a number of basic checks are presented, comparing the model of ELENA built with SIMPA to the MAD-X [18] model. In the frequency analysis section, the effect of fringe fields and the magnetic system of the electron cooler is investigated. The dynamic aperture section quantifies the effect of resonance lines identified as relevant by the frequency analysis. In the conclusions we suggest a few possible improvements for the performance of ELENA.

II. THE TRACKING ALGORITHM

A. A short summary

We recommend reading the previous paper [5] to understand the algorithm in detail, as only a summary is provided here.

Symplectic integrators keep the conserved quantities bounded, but cannot cure the errors coming from the representation of the fields. These are two separate sources of errors. It is crucial to have a physically valid representation of the fields. When an electromagnetic field not satisfying the Maxwell equations is given to a symplectic integrator the integration will be still symplectic, but the result will describe something other than the true physical behavior of the system. A common example is the energy drift due to spurious sources appearing in 3D grid methods [7].

The potentials describing the fields must satisfy the Laplace equation and should be continuous everywhere in the beam region to obtain accurate results in long-term trackings. To comply with these conditions, the potentials are expressed analytically in terms of their sources. We use hypothetical magnetic and electric monopoles and current loops as sources. These sources are placed outside of the volume of interest, at some distance from the boundary, and their strength is set such that they reproduce the normal component of the magnetic or electric field at the boundary by solving a system of linear equations.

The normal component of a magnetic or electric field determines completely the field inside a closed simply-connected volume. For a multiply-connected volume, this is not always true. If the integral of the magnetic field along a closed curve is not zero, then magnetic monopoles alone cannot reproduce the magnetic field. One example of this situation is the solenoid magnet. In these cases, the nonconservative part of the vector potential must be

provided by one or more current loops. The word non-conservative should be understood here in a mathematical sense, the vector field having a nonzero integral along a closed loop, that is $\oint \mathbf{v} \cdot d\mathbf{r} \neq 0$.

After the potentials are reproduced at the boundary by the sources, they can be evaluated anywhere inside the volume analytically. The potentials satisfy the Laplace equation and they are continuous. However, this method is too slow to be practical. We will refer to this part of the algorithm as slow evaluation. We need a much faster calculation of the potentials to be of practical use.

Several orders of magnitude improvement can be achieved by using a local description of the potentials. Spherical harmonics scaled appropriately are called solid harmonics. Regular solid harmonics are the canonical representation for harmonic functions inside a sphere [19]. A key characteristic of the algorithm is the description of vector and scalar potentials by solid harmonics inside a set of overlapping spheres covering the volume of interest.

The potentials satisfy exactly the Laplace equation inside the spheres. The discontinuity between the spheres decrease exponentially with the degree of solid harmonics expansion and can be easily kept close to machine precision. The representation of the potentials in terms of solid harmonics is optimal in terms of memory and allows fast evaluation. We will refer to this part of the algorithm as fast evaluation. The speed of the evaluation is strongly dependent on the degree of the solid harmonics expansion. This property is the base for the speed improvements we will describe later.

One might ask what is the point of having a slow evaluation when we have a fast one. To calculate the solid harmonics coefficients for each ball in the fast evaluation we need the values of the scalar and vector potentials. These potentials must be continuous and must satisfy the Laplace equation. This is the role of slow evaluation. CAD software or measurements can not provide the potentials satisfying these conditions. We use the field values obtained from those at the boundary of the volume of interest to determine the strength of the monopoles and current loops of the slow evaluation. The workflow starts with modeling individual magnets with CAD software or measurement data, then the strengths of the monopoles and current loops are calculated for each type of magnet. The next step is to assemble these sources according to the lattice of the machine. Finally, the solid harmonics coefficients of each sphere covering the volume of interest are calculated, providing a field map for the fast evaluation.

B. Extension to general domains

In [5] we showcased the algorithm on a simple accelerator. The beam region had a circular shape everywhere with a constant 3 cm radius all along the reference orbit. This was covered by a single row of overlapping spheres with a 4 cm radius. This solution to cover the beam region is sufficient to demonstrate the method, but real accelerators have a more complex aperture in general. Apart from that, our algorithm aims to apply to a wide range of problems, not only in the domain of particle accelerators.

The usual way to describe a general volume is to give its boundary surface by triangulation. A common format to describe a triangulated surface is the standard tessellation language abbreviated as STL. The boundary of the volume of interest is described by an STL file in SIMPA. There are many CAD software able to produce STL files, however, to construct the aperture for an accelerator using CAD software is not a trivial task. There is a utility implemented in SIMPA facilitating the construction of the STL file. It takes the design orbit and a list of transverse profiles as the input and extrudes those profiles along the design orbit. For each transverse profile, the user must specify the longitudinal position where the extrusion starts and ends along the orbit. The transverse profiles are described as polygons. The number of segments for each polygon is restricted to the same value, making the implementation of the triangulation straightforward.

Once the boundary of the beam region is specified by an STL file the next step is to fill it with overlapping spheres, such that the entire beam region is covered without gaps. There are constraints concerning the size of the spheres. The slow evaluation calculates the potentials from sources, which are placed at some distance from the boundary. The spheres should be small enough to not overlap with the sources. The solid harmonics expansion works only if there are no sources inside the spheres.

It is practical to place the center of the spheres onto a regular lattice. This simplifies the construction of the cover as well as finding the spheres when the potentials are evaluated. There are many possible lattices, but only the hexagonal close-packed (HCP) and the face-centered cubic (FCC) lattices have the optimal packing density, about 0.74. We chose the HCP lattice. The coordinates of the sphere centers on an infinite HCP lattice can be obtained by three simple expressions. We keep only those spheres from the infinite lattice which are necessary to cover the beam region.

The first step is to find the indices of one initial sphere on the lattice, which has its center inside the beam region. Then a 3D version of an algorithm from computer graphics, called flood-fill, is applied to collect all sphere centers inside the boundary. In essence, the flood-fill algorithm checks neighboring lattice locations recursively, if they are inside or outside the boundary. After this step, we have a list of spheres which covers most of the beam region, but not all. Near the boundary, there can be gaps.

A second step is necessary to cover the beam region fully. In the second phase, all spheres kept in the first step are checked in the following way. If the sphere in question has missing neighbors in the lattice, then we take the half-line from the sphere center to the center of the missing
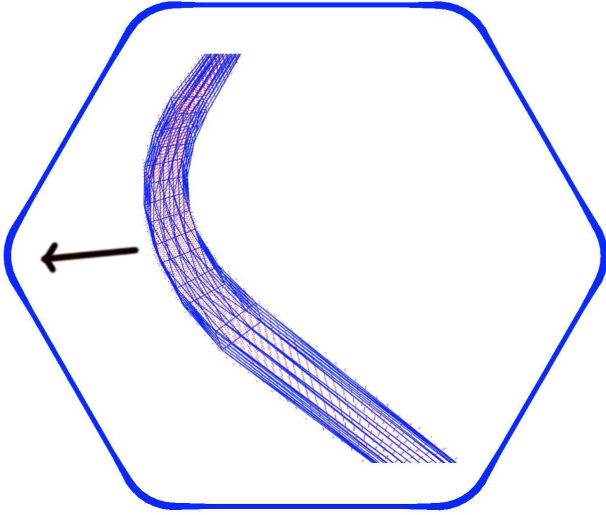
FIG. 2. Triangulated boundary of the ELENA beam region. The interior plot is an expanded 3D view of an arc. The red dots are the centers of the covering spheres located on an HCP lattice.

neighbor and calculate its intersection point with the sphere in question. If this point is inside the boundary of the beam region, then we add the missing sphere to the cover. We do this for all missing neighbors, hence eliminating gaps. We should note this second step can be applied only if the distance of sources providing the potentials are further from the boundary than the diameter of the covering spheres. In the present study that is the case. Otherwise, some spheres would have sources inside, which is not allowed. A more sophisticated second step of the covering process could be applied by removing the restriction to put sphere centers only to the lattice points and allowing variable sphere sizes in the second step. This is a nontrivial problem in computational geometry, but certainly has a satisfactory solution. Figure 2 shows the ELENA beam region and the centers of the covering spheres. The radius of the spheres was 1 cm in this study and 75951 of them were needed to cover the aperture.

C. Improvements in speed and accuracy

Covering the volume of interest as described above makes the method applicable to a much wider range of problems, but this is not the only advantage. In general, the error of the spherical harmonic approximation on the surface of a sphere is less than the first term omitted from the infinite sum of spherical harmonics times a constant [19]. Let us inspect Eq. (9) in [5], which is restated here for convenience.

$$f(r, \theta, \phi) = \frac{1}{R} \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} r^\ell c_{\ell m} \tilde{P}_\ell^m(\cos\theta) \Phi(\phi; m), \quad (1)$$

where $\Phi(\phi; m)$ given by

$$\Phi(\phi; m) = \begin{cases} \sqrt{2}\sin(|m|\phi) & \text{if } m < 0 \\ 1 & \text{if } m = 0 \\ \sqrt{2}\cos(m\phi) & \text{if } m > 0 \end{cases} \quad (2)$$

$\tilde{P}_\ell^m(\cos\theta)$ is the orthonormalized associated Legendre polynomial, $c_{\ell m}$ is a coefficient, R is the radius of the sphere and $r = |\mathbf{r_e} - \mathbf{r_c}|/R$ is the scaled distance between the sphere center $\mathbf{r_c}$ and the evaluation point $\mathbf{r_e}$. We can set $r = R$ and calculate the maximum error of the solid harmonics approximation as $\epsilon R^{\ell_{\max}}$, where ϵ is the error due to the first omitted term in the sum with degree $\ell_{\max} + 1$. Because of the presence of the term r^ℓ, the biggest possible error in the approximation strongly depends both on R and $\ell_{\max}$.

For instance, in [5], we used a single row cover with $R = 4$ cm and $\ell_{\max} = 38$. In this study $R = 1$ cm. If we had used the same $\ell_{\max} = 38$ with $R = 1$ cm spheres, the maximum error of the approximation $\epsilon R^{\ell_{\max}}$ would have been a factor 4^{38} smaller, assuming infinite precision arithmetic and the two spheres having the same center. This is a big number, so we can reduce $\ell_{\max}$ and the error of the solid harmonics approximation still stays close to machine precision.

In this study we used $\ell_{\max} = 12$ uniformly for all spheres. This speeds up the evaluation of the potentials and their derivatives significantly, because the number of operations to calculate the sum in Eq. (1) is proportional to $\ell_{\max}^2$. On Intel i7-6700 CPU at 3.40 GHz, using a single core with hyper-threading enabled, we got 2.7×10^5 evaluations per second of the vector potential and its derivatives with $\ell_{\max} = 12$. This should be compared against 5.8×10^4 evaluations with the single row cover and $\ell_{\max} = 38$, a factor 4.66 improvement.

Further speed improvement can be achieved by changing $\ell_{\max}$ between the spheres to keep the error of the approximation below a prescribed value. There is also ample space for optimization by improving memory locality, better use of vector instructions of the CPU, and parallelization. This will be implemented in the future.

There is one downside of the reduced sphere diameter, that is the increased memory use. The number of spheres needed to cover a given volume is proportional to $1/R^3$. The number of coefficients in a solid harmonics approximation is proportional to $\ell_{\max}^2$, which in turn depends on R. Together, the number of coefficients needed to describe a potential in a given volume with some prescribed accuracy is $n = c\ell_{\max}^2/R^3$, where c is a constant. The constant c depends on the volume to be covered. In the present study the size of the field map is 448 MB. Using a single row cover with $\ell_{\max} = 38$ it is only 29 MB. Memory is traded for speed.

III. PREPARATIONS

A. Preparing the field maps

As a first step, the field of each type of magnet in ELENA has to be expressed as a collection of magnetic monopoles or a combination of magnetic monopoles and current loops. This provides continuous and analytic potential everywhere in the beam region. To do so, the magnetic field values have been obtained from the CAD software OPERA for each magnet at specific points on a surface surrounding the beam region. This surface is close to the poles of the magnet. The magnetic monopoles are placed outside of the surface at 2.5 cm distance. Then a system of linear equation was solved to find the strengths of the monopoles for each magnet type. The relative precision of the reproduction of the magnetic field from the vector potential is typically between 10^{-3} and 10^{-4} in this study. The precision can be made better by increasing the number of monopoles.

Once each type of magnet in the machine has been expressed as a collection of monopoles and current loops, the next step is to assemble the ELENA ring from these collections. This has been done by rotating and translating the collections of sources to the correct place. The entry and exit coordinates of the magnets have been obtained from the MAD-X model with the survey command.

After the machine is assembled, a particle can be tracked to check if the procedure used so far has been correct. This tracking is slow because the vector potential of all sources has to be evaluated at each time step. Nevertheless, it is fast enough to track a few 10's of turns to calculate the closed orbit and set the values of the orbit correctors, which is necessary when the electron cooler is included.

The magnets of the ELENA ring have been organized into magnet groups. Usually, a magnet group consists of magnets connected to the same power supply. For each magnet group, a field map was produced. These field maps consist of a large number of overlapping spheres, each with some number of coefficients determined by the maximum degree of solid harmonic expansion as described earlier. The location and size of the spheres are calculated once and it is common for all field maps. This makes it easy to scale field maps individually and sum them into a single map before tracking.

B. Basic checks

It is reassuring to implement some checks and compare the results to an established tool when this is possible. Since there is already a MADX-PTC model available for ELENA, we use this model as a reference for comparison. There are several global parameters of the machine which can serve as cross-checks, such as the tunes and the optical functions. We compared the optical β and dispersion functions to those obtained with MAD-X at the same currents for the quadrupoles. In this comparison, the

electron cooler magnets and the solenoid compensators were turned off in both models as well as all orbit corrector magnets.

In our model, the optical functions are obtained from tracking. This might look more complicated than the usual linear optics model, but it is necessary because we deal with arbitrary 3D fields. To obtain the Twiss functions, we tracked a $\bar{p}$ with nominal energy while saving the phase space variables at each time step. The symplectic integration method we implemented according to [17] uses canonical variables, but these were converted to positions and angles. The procedure results in many Poincaré maps along the orbit. We then fitted an ellipse to the phase space variables at each longitudinal position. From the parameters of the ellipse, the Twiss functions were calculated by well-known formulas. To get the dispersion, another particle with a different dp/p was tracked and the center of the fitted ellipses was compared to the $dp/p = 0$ case. The β and D functions are plotted in Fig. 3. The agreement with MAD-X is remarkable if one considers the completely different calculation of the Twiss functions.

The tunes were also compared. We chose the working point $Q_h = 2.454$ and $Q_v = 1.416$. The reason for this choice will be clear later in light of the frequency analysis results. The tunes calculated by MAD-X with the PTC [20] Twiss command are reasonably close, $Q_h = 2.479$, and $Q_v = 1.42$.

A $\bar{p}$ with $dp/p = 0.001$ was tracked and compared to the $\bar{p}$ with nominal energy to calculate the linear chromaticity. The tunes for both $\bar{p}$ were calculated from the tracking results with a very precise frequency analysis tool called PyNAFF [21]. The chromaticities we got are $\xi_h = -4.72$ and $\xi_v = -0.15$, which is significantly different from the results obtained with MADX-PTC, $\xi_h = -3.3$ and $\xi_v = -1.81$. One possible explanation for this discrepancy is that the end fields in MADX-PTC are taken into account
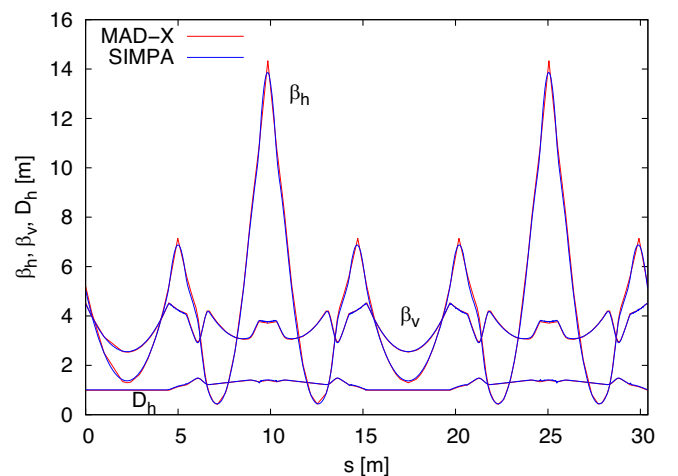

FIG. 3. Comparison of the Twiss functions β_h, β_v and D_h calculated by SIMPA and MAD-X at the working point $Q_h = 2.454$ and $Q_v = 1.416$. The functions are hardly distinguishable on the plot.

by the FINT and HGAP parameters [18,22] for the bending magnets, while our algorithm handles the end fields nearly exactly for any element. We have not made measurements during commissioning with similar conditions in ELENA.

IV. FREQUENCY ANALYSIS

A. Motivation

Small machines have special challenges, which can be neglected in large rings. One of these is the increased influence of fringe fields on the beam dynamics [2]. ELENA is a ring with 30.4 meters circumference. It is mandatory to take into account the effect of fringe fields already during the design of machines of this size. The ELENA MADX-PTC model includes the fringe fields for the main bending magnets, which is one of the main sources of nonlinearities. Another important source of nonlinearities is the magnetic system of the electron cooler, but it is not included in the MADX-PTC model. Our model includes all fringe fields and also the magnetic field of the electron cooler with high fidelity.

B. The method

According to the KAM theory, invariant tori of a conservative dynamical system are stable under small perturbations. These tori describe quasiperiodic orbits in phase space. The motion of the system can be described by action-angle variables. On invariant tori the action variables are constant and the angle variables have fixed irrational and Diophantine frequencies. When the frequencies are rational numbers, they are related to each other by integers and resonance occurs. The condition for resonance in a particle accelerator in 4D phase space is given as

$$nQ_h + mQ_v = p, \tag{3}$$

where Q_h, Q_v are the horizontal and vertical tunes and n, m, p are integers, not all zero. A more detailed accessible explanation can be found in [23] for instance. Although the frequencies where a resonance occurs are known *a priori*, their strength or significance is not known. One can track particles to investigate the strength of the resonances. Doing so for a large number of turns is time-consuming. Fortunately, there are some early indicators to detect chaotic behavior. One early indicator is the frequency shift.

An early use of frequency analysis for conservative Hamiltonian systems by Laskar [24] was to investigate the stability of the solar system. Later it was applied to other domains among them accelerator physics [25–27]. For a comprehensive description of the method see [28]. The precision of Laskar's NAFF [29] algorithm is the feature that makes it interesting. While the frequency error of the standard FFT is proportional to $1/N$, the error of NAFF is proportional to $1/N^4$, where N is the number of samples.

This allows us to determine the characteristic frequencies of the motion with a small number of turns.

The first characteristic frequencies for each plane of the motion are the tunes. The drift in the tunes can serve as an early indicator of the long-term stability of the motion [30]. For initial conditions corresponding to chaotic trajectories, the frequency can only be defined for a given time interval. Calculating this local frequency for two consecutive time intervals and taking their difference we can calculate tune shifts $\Delta Q_{h,v}$. For a coasting beam, the figure of merit can be defined as:

$$D = \log_{10}\left(\sqrt{\Delta Q_h^2 + \Delta Q_v^2}\right) \tag{4}$$

D calculated with this definition can be used to identify stable areas in the tune diagram.

To identify promising working points, we scanned a large range of tunes in the horizontal and vertical planes and for each initial condition the figure of merit D was calculated. The ELENA machine has three quadrupole families, which give three parameters for the tune scan. There is also a constraint for the dispersion in the electron cooler region, which is fixed to 1 meter. This constraint selects a 2D plane in the 3D space of possible quadrupole currents. The two-dimensional tune space was scanned between 2.2–2.52 and 1.2–1.5 horizontally and vertically by changing the currents in the three quadrupole families. The settings of quadrupole currents for the estimated tunes were calculated by a second-order polynomial. This polynomial is identical to the one used in the ELENA control system. The estimated tunes are a bit different from the tunes obtained from the tracking results by PyNAFF, which we call "numerically measured tune." This is because the settings of quadrupole currents for the estimated tunes from the second-order polynomial is only an approximation based on linear optics.

The ranges of tunes in both directions are large, and the optics can change significantly during the scan. The initial conditions have to be varied to compensate for the optics change and keep the emittances approximately constant. The emittances were kept 30 π mm mrad in both planes. This particular value for the emittance was chosen to be large enough to see clearly the nonlinear effects due to the fringe fields and the electron cooler, but without having too much losses due to the physical aperture of the machine. In this paper emittances and acceptances are defined as the area of the ellipse corresponding to the single particle Courant-Snyder invariant, unless stated otherwise.

The tunes were scanned in 160 steps in both directions giving 25600 initial conditions. Each particle was tracked for 300 turns and the phase space variables were saved at a single longitudinal position for each turn. The tracking result was post-processed with PyNAFF, a Python implementation of the NAFF algorithm [21]. To calculate $\Delta Q_{h,v}$, we split the 300 turns into two sets. Then the figure of merit

D in Eq. (4) was calculated for each point in the diagram and plotted as colors.

C. Results

Figure 4 shows the figure of merit D for the bare ELENA machine with bendings and quadrupoles only, plotted against the estimated tunes as colors. The term 'estimated tune' means the input for the polynomial, mapping the desired tunes to quadrupole currents, which is slightly different from the actual tunes. The red dots indicate the initial conditions which resulted in particle loss in less than 300 turns.

Instead of the estimated tunes, we can plot D against the tunes calculated by pyNAFF as shown in Fig. 5. This gives a precise value for the tunes, but this tune can not be calculated for the lost particles, that is why Fig. 5 has no red dots inside. Figures 4 and 5 are somewhat complementary.

On both plots, many higher-order resonance lines can be seen. It should be emphasized, these resonance lines are present also in a perfectly manufactured, built, and aligned machine. We have not put any imperfections into the model. All the resonance lines are the results of the fringe fields of the bending magnets and the quadrupoles. They are the direct consequence of the geometry of the magnets.

The next step was to repeat the exercise with the magnetic elements of the electron cooler included. Those are the main solenoid coil, two toroid coils at the ends, two solenoid compensator magnets, and several orbit correctors including the two "Kyoto type." We did all tracking at the lowest energy of 100 keV, because the effect of the electron cooler is most pronounced at this energy. The cooler has fixed currents for all coils during the deceleration cycle.
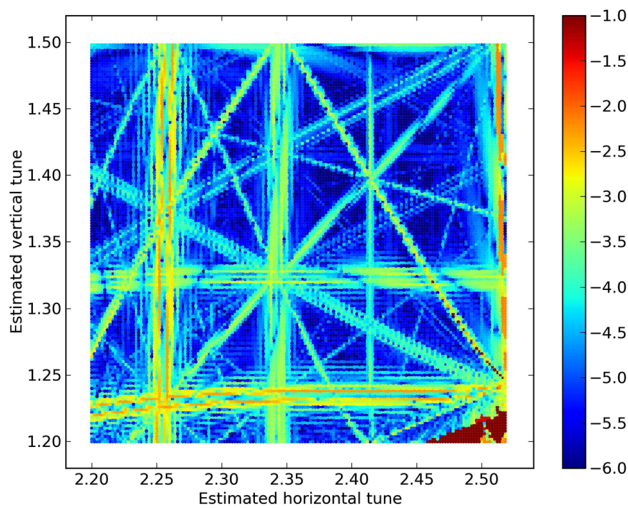

FIG. 5. D plotted as colors against the numerically measured tunes for the bare ELENA machine, consisting of only the main bendings and the three quadrupole families.

The step size was set to 1 cm while tracking the bare machine. With the electron cooler included, it was reduced to 0.5 cm. This was needed because the cooler introduced some spatially fast-changing magnetic fields. We used an explicit second-order integrator described in [17]. Figure 6 shows D against the estimated tunes and Fig. 7 against the numerically measured tunes. It is immediately apparent that the magnets of the electron cooler have a significant effect on the beam dynamics. Many resonance lines became stronger and wider and particle losses are more frequent at the strongest resonance lines.

It is natural to ask, how much space is available for the beam in these diagrams. To answer this question we sampled 1000 particles from a Gaussian distribution with

FIG. 4. D plotted as colors against the estimated tunes for the bare ELENA machine, consisting of only the main bendings and the three quadrupole families. The red dots correspond to lost particles.
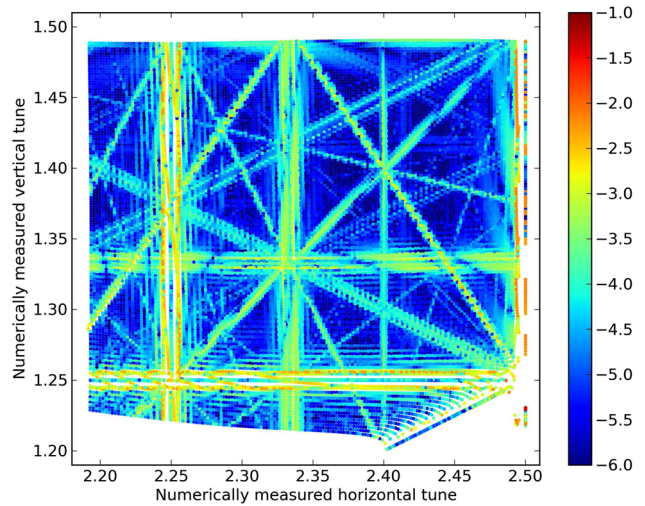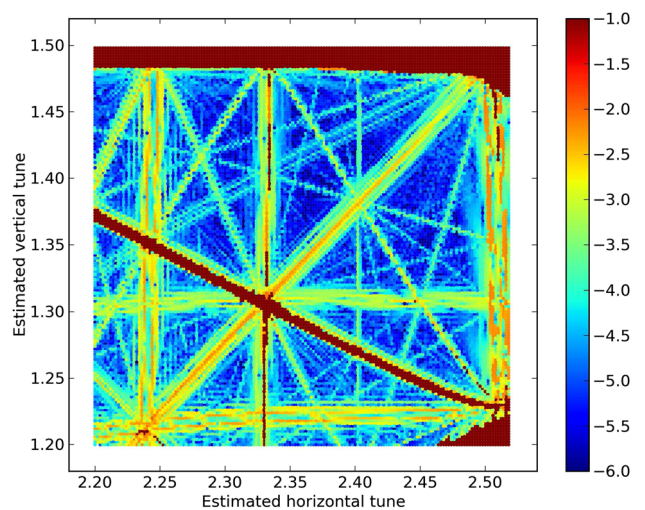

FIG. 6. D plotted as colors against the estimated tunes for the ELENA machine with electron cooler. The red dots correspond to lost particles.
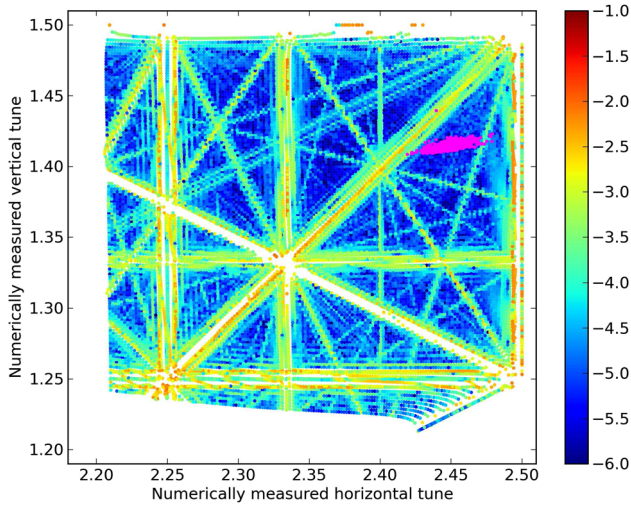
FIG. 7. *D* plotted as colors against the numerically measured tunes for the ELENA machine with electron cooler. The magenta dots in this plot correspond to the tunes of particles which survived 10^4 turns in a bunch filling the acceptance, without space charge included in the model.

2σ RMS emittances 75 π mm mrad in both planes, which is the acceptance of the machine and $dp/p = \pm 9 \times 10^{-3}$ which is about the momentum acceptance given as 2σ RMS. These particles were tracked with the machine set to the working point $Q_h = 2.455$ and $Q_v = 1.415$. This working point gives the most space free of strong resonance lines. Then the tunes were calculated for each surviving particle with PyNAFF. The tune footprint is inserted into Fig. 7.

Several effects contribute to the tune footprint. The chromaticity is an important one, specially the horizontal. With the electron cooler included $\xi_h = -4.4$, $\xi_v = -0.55$. Another contribution comes from the betatron amplitude dependence of the tunes. The third and most challenging one to deal with is the space charge. The effect of space charge is not investigated in this work, but is clearly an important factor to consider. According to [6], the incoherent tune shift can be -0.1 for a short time before extraction when the beam is bunched. For a coasting beam it is much less, about -0.01, but for a much longer time in the order of seconds. Even in the case of coasting beam the contribution of the space charge to the tune footprint is significant. Therefore it is important to choose a working point which have a sufficiently large space around it in the tune diagram without strong resonance lines, which is the working point $Q_h = 2.455$ and $Q_v = 1.415$.

V. DYNAMIC APERTURE

Frequency analysis allows one to compare and select the promising working points by giving estimates of the relative strength of resonances in the tune diagram. The next step is to analyze the effect of these resonances on the beam. The concept of dynamic aperture quantifies the available stable phase space volume for the beam.

A. The method

The stability domain is a connected region in phase space of initial conditions that remains bounded after N turns [31].

The stability region is the area inside the last simply-connected invariant curve in phase space. In 4D of phase space, the invariant curves do not always separate a connected region, but in most practical cases it exists. We will see later in the results that even if the working point is on a resonance line a connected region can be identified.

In 4D of phase space, the volume of the stability domain can be calculated by the integral

$$\int \int \int \int \chi(x, p_x, y, p_y) \, dx \, dp_x \, dy \, dp_y, \qquad (5)$$

where $\chi(x, p_x, y, p_y)$ is a function of initial conditions, with value one if the particle has survived N turns, or zero if it was lost.

The dynamic aperture is defined as the radius of the hyper sphere having the same volume in phase space as the stability domain. In practical calculations the phase space coordinates (x, p_x, y, p_y) of Eq. (5) are expressed in terms of polar coordinates α, r in physical space and the phase space angles v_1, v_2, then the integral approximated as a sum over these coordinates.

The direct calculation of Eq. (5) by summation over all four variables is very CPU time demanding. Instead we applied the method described in [31,32] as "integration over the dynamics." The idea is to replace the integration over the phase space angles v_1, v_2 by averaging over the dynamics. While particles take turns, they sweep the full range of phase space angles v_1, v_2, and this information is available.

The procedure is the following. With fixed initial phase space angles $\bar{v}_1$, $\bar{v}_2$, the polar angle α_k is scanned in K steps. For each α_k the radius of the last stable orbit $r(\alpha_k, \bar{v}_1, \bar{v}_2)$ is sought by increasing the polar radius variable in J steps. Then the N turns of the last stable orbit for each α_k is distributed into an angular grid of size $M \times M$, such that each grid element formed by the phase space angle intervals Δv_1, Δv_2, should contain at least one iterate from the orbit of $r(\alpha_k, \bar{v}_1, \bar{v}_2)$. Then the dynamic aperture is calculated with the formula

$$r_d = \left[\frac{\pi}{2KM^2} \sum_{m_1,m_2=1}^{M} \sum_{k=1}^{K} r_{m_1,m_2}(\alpha_k, \bar{v}_1, \bar{v}_2)^4 \sin(2\alpha_k) \right]^{\frac{1}{4}}, \quad (6)$$

as given in [31].

We found that M could not be always increased above a rather low value without having at least one empty bin in the $M \times M$ grid. This is due to the very inhomogeneous distribution of phase space angles v_1, v_2, which is a known phenomenon [31,33]. A low value for M leads to low accuracy.

There is another way to calculate the dynamic aperture based on normal forms given to the first-order in [31] as

$$r_{NF} = \left[\frac{\pi}{2K} \sum_{k=1}^{K} \{\rho_1 + \rho_2\}^2 \sin(2\alpha_k) \right]^{\frac{1}{4}}, \qquad (7)$$

where $\rho_1(\alpha_k, \bar{v}_1, \bar{v}_2)$ and $\rho_2(\alpha_k, \bar{v}_1, \bar{v}_2)$ are the nonlinear invariants, which are the same as the Courant-Snyder invariants. Since there is already an implementation of ellipse fit to the phase space variables in SIMPA, the values of $\rho_1(\alpha_k, \bar{v}_1, \bar{v}_2)$ and $\rho_2(\alpha_k, \bar{v}_1, \bar{v}_2)$ are easily obtained. It should be noted however, this sum can lead also to non-negligible errors if the working point is on a strong resonance line and the form of the phase space plot is significantly different from an ellipse. Nevertheless, it is still the best first-order approximation of the dynamic aperture. For this reason we calculated with both Eq. (6) and Eq. (7) as a cross check.

In our numerical experiments, the closed orbit was different from the design orbit. The electron cooler has introduced an orbit distortion, which was corrected. However, the correction was not perfect. The maximum orbit deviation was less than 3 millimeters horizontally and less than 1 millimeter vertically. It could have been reduced further, but doing so for each point we investigated is rather tedious. Also, in reality, the orbit of ELENA during the commissioning was bigger than that.

B. Results

The dynamic aperture of six working points has been evaluated to compare resonance conditions with a non-resonant case. The selected set of tunes is indicated on Fig. 8. These resonance lines were selected because they are close to the region where the working point of ELENA was set during the commissioning. The first working point was chosen to be far from strong resonance lines at $Q_h = 2.455$, $Q_v = 1.415$. The other five points were placed on resonance lines with various orders. The dynamic apertures calculated with Eq. (6) and Eq. (7) are displayed in Table I. All the results were calculated with $N = 10^4$ turns, $K = 20$ polar steps, with $\alpha_k \in [0, \pi/2]$, radial steps $J = 20$ and the biggest possible M, which is also indicated in Table I. The energy of the particles was set to 100 keV. We also calculated the errors of the dynamic aperture values in Table I with Eq. (10) given in [34], it is ± 0.0011 meter.

Most of the dynamic aperture values in Table I are close to 0.011. At first sight one is tempted to conclude that the
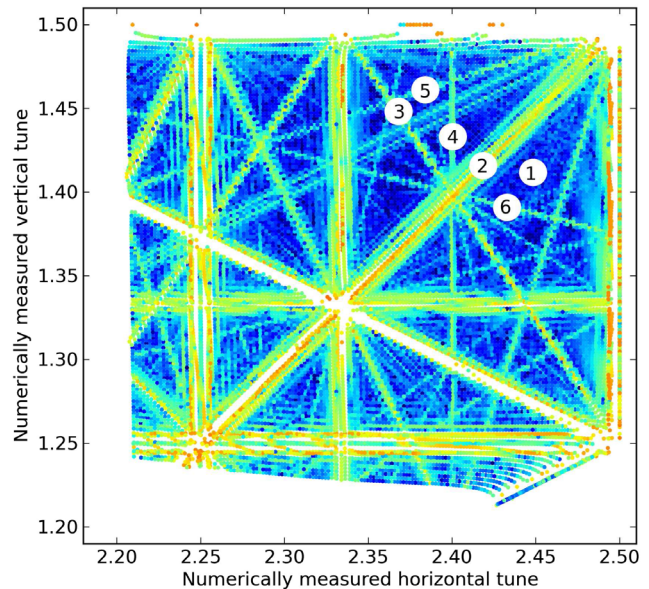

FIG. 8.   The numbered white points in the tune diagram indicate the tunes where the dynamic aperture was calculated.

resonance lines of order 4 or 5 are already weak enough to not reduce further the physical aperture of the machine. The dynamic aperture however is defined as the radius of a four dimensional hyper sphere and small differences in the radius can hide significant phase space volume reductions at specific angles of α_k.

For this reason we found it informative to give also a more detailed description by plotting the stability diagrams in Fig. 9 for each point in Fig. 8, which survived $N = 10^4$ turns. In the stability diagrams, instead of the usual convention of using the physical beam dimensions in units of sigmas of the beam distribution on the axes, we used the single-particle emittances. In ELENA, the beam size depends on the performance and duration of the electron cooler. We repeated the same exercise with particles having a momentum deviation $dp/p = 2 \times 10^{-3}$, which is the maximum momentum spread measured during the commissioning. The quadrupole currents were adjusted for each

TABLE I.   Dynamic apertures for the six points in Fig. 8 calculated with averaging over the dynamics (r_d) and with normal forms (r_{NF}). M is the angular grid size for r_d calculation. For the tunes falling onto a resonance line, the equation of the resonance is also given.

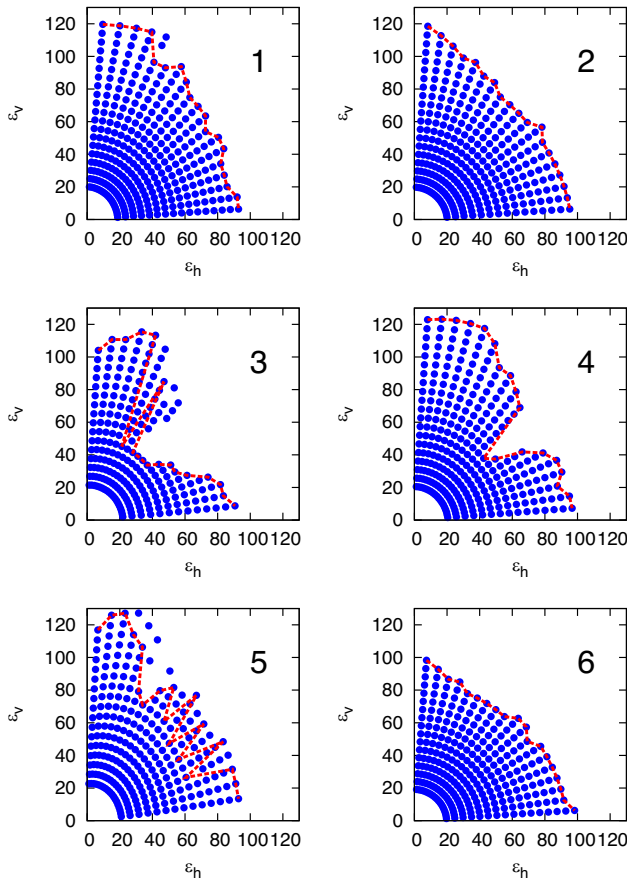| Point	r_d [m]	M	r_{NF} [m]	Resonance condition
1	0.0115	14	0.0115	NA
2	0.012	8	0.0112	$Q_h - Q_v = 1$
3	0.01	6	0.01	$3Q_h + 2Q_v = 10$
4	0.0117	3	0.0114	$5Q_h = 12$
5	0.0113	9	0.011	$-Q_h + 3Q_v = 2$
6	0.0113	8	0.0108	$Q_h + 4Q_v = 8$

FIG. 9. Initial emittances plotted in units of π mm mrad for particle with $dp/p = 0$. The horizontal axis corresponds to the horizontal initial emittance and the vertical axis to the initial vertical emittance. Only those initial conditions are plotted which survived $N = 10^4$ turns. The numbers in the upper right corners correspond to the numbering in Fig. 8. The lines indicate the last connected initial emittances for each angle α_k.

point to have the same tunes as for the case $dp/p = 0$. The results are plotted in Fig. 10.

In Fig. 9, 10 it is apparent that some of the resonance lines make the dynamic aperture significantly smaller than the physical aperture. The transverse acceptance of ELENA is determined by the size of vacuum chambers and the optics.

There are a number of interesting features displayed in these figures, some of them surprising. The coupling resonance $Q_h - Q_v = 1$ indicated by point two is rather wide and with high diffusion rate, but it does not significantly reduce the dynamic aperture. The smallest dynamic aperture was obtained at point three, due to a 5th order resonance and not a lower-order resonance line. We did not include first, second, and third-order lines, those should obviously be avoided during operation. Higher-order resonance lines however can not always be evaded. The incoherent space charge tune shift in ELENA at the lowest energy is significant. Its estimated value for a bunched
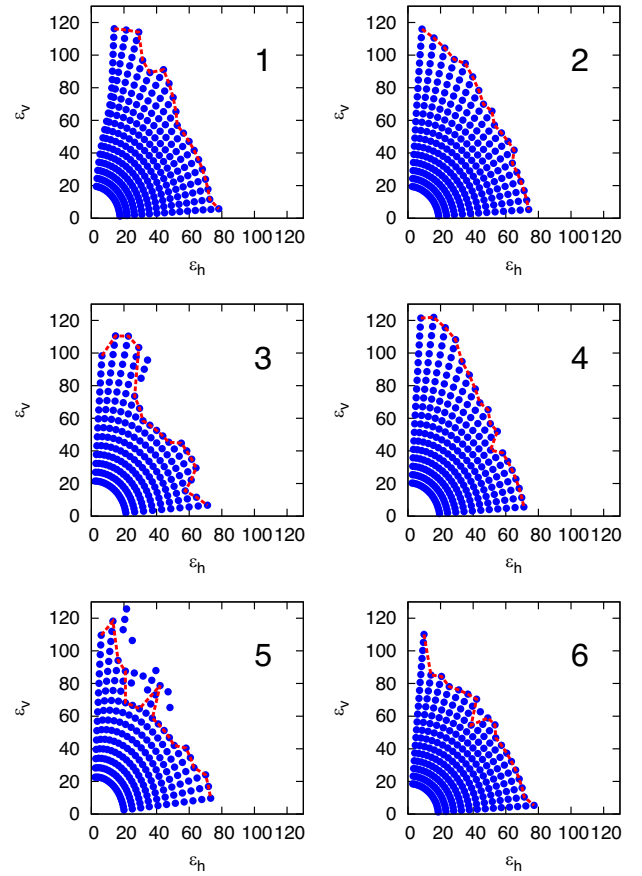

FIG. 10. Initial emittances plotted in units of π mm mrad for particle with $dp/p = 2 \times 10^{-3}$.

beam at arrival to 100 keV is $\Delta Q = -0.03$ [6]. The incoherent space charge tune shift contributes to the tune footprint of the beam and some of the resonances shown in Fig. 8 are unavoidable. Therefore it is important to choose a working point which gives the most space free of harmful resonance lines. Point one in Fig. 8 is a good choice, because the surrounding resonance lines $Q_h - Q_v = 1$ (on point two) and $Q_h + 4Q_v = 8$ (on point six) are rather forgiving according to Fig. 9, 10.

The initial emittances for the test particle to be tracked were calculated by a linear model, that is the Courant-Snyder theory. It assumes that the values of the phase space variables lie on an ellipse. Under resonance conditions, the phase space ellipse is distorted. Although the angle α_k was changed linearly in equal steps, in Fig. 9 case 5 for instance, shows distortions. This is because the phase space plot deviates significantly from an ellipse and the requested emittance is different from the actual initial emittance, which is plotted.

The dynamic aperture is defined as a function of turns. It is known to decrease with the inverse of the logarithm of the number of turns [34–36]. Indeed, it is easy to find initial conditions which leads to particle loss after more than 10^4 turns, an example is given in Fig. 11.
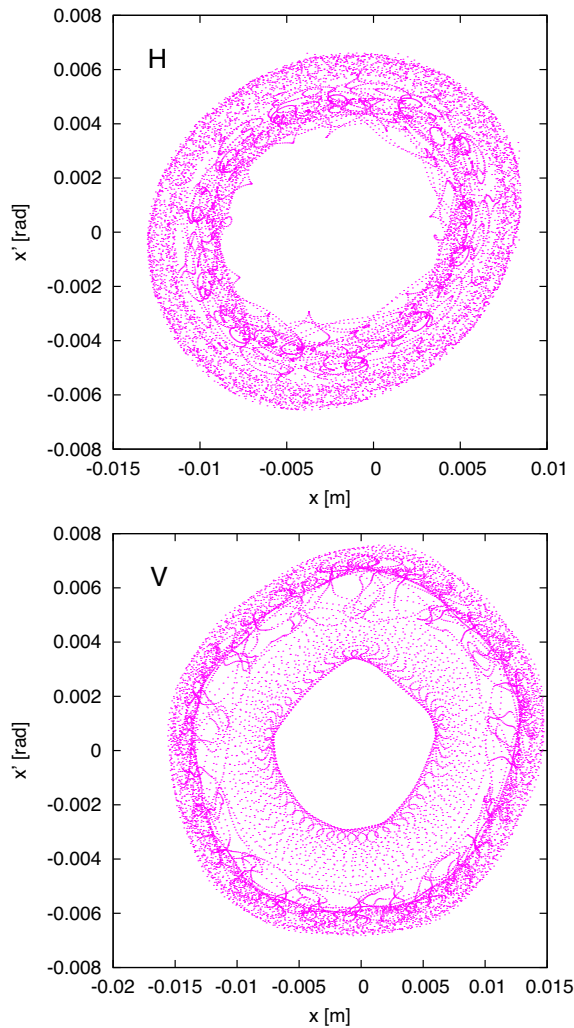
FIG. 11. Phase space plots of a particle near a 4th order resonance, lost after 14211 turns.

VI. CONCLUSIONS

We have applied the SIMPA code for long-term symplectic charged particle tracking in arbitrary static electromagnetic fields on the ELENA machine. The algorithm described in a previous paper has been extended to handle complex beam region geometries. A short description of the enhanced algorithm was given. The improvements not only raised its generality, but also increased the speed by a factor 4.66.

The frequency and dynamic aperture analysis identified a number of 4th and 5th order resonance lines in the tune diagram strong enough to reduce the dynamic aperture below the physical one. What made the frequency and dynamic aperture analyses different, is the fact that we have not introduced any multipole error into our model apart from the multipole components due to the geometry of the magnets, which are inevitable. All the resonances seen in the frequency analysis are direct consequences of the geometry of the fields, even if the magnetic elements are manufactured perfectly.

We showed in the frequency analysis section the effect of the magnetic system of the electron cooler on the beam dynamics by comparing the two cases, with and without electron cooler. The electron cooler introduced non-negligible magnetic perturbations, strengthening many resonance lines. To our best knowledge there was no similar study done before with electron coolers.

A. Suggestions to improve the performance

The commissioning of the ELENA ring was finished in November 2018, at the start of the 2.5 years long CERN wide shutdown. The deceleration efficiency was estimated between 40–50%, somewhat below the 60% design value. Most of the losses occurred on the ramp just before the lowest energy plateau. This is not surprising, since both the incoherent tune shift due to space charge and the intrabeam scattering has the biggest effects at low energy. Also the perturbation introduced by the electron cooler is the strongest at the lowest energy.

During the commissioning two working points were explored extensively, $Q_h = 2.42$, $Q_v = 1.45$, and $Q_h = 2.38$, $Q_v = 1.44$. These working points are not far from the 4th and 5th order resonances number 3, 4, 5 in Fig. 8 and these resonance lines are not harmless according to Fig. 9, 10. The tune footprint of the beam certainly overlaps with some of the resonance lines. The incoherent tune shift of a bunched beam due to space charge at arrival to 100 keV energy is on the order of −0.03, depending on the beam intensity and beam size. This is not a negligible value, but we believe it is not big enough to drastically change the results of our analysis. One supporting evidence for this assumption, is that we have not observed significant dependence of the deceleration efficiency on the beam intensity during the commissioning. After cooling at 100 keV energy, the incoherent tune shift is in the order of −0.1, but only for the duration of the final part of the bunching before extraction, about 10 milliseconds. No significant losses were observed related to the bunching process during the commissioning. Based on the frequency and dynamic aperture analysis, we suggest to explore working point number one in Fig. 8, as it gives the most available space without harmful resonances.

The tune footprint of the beam with maximum emittances in all planes as shown in Fig. 7 is mostly coming from the horizontal chromaticity which is $\xi_h = -4.4$ at 100 keV energy with electron cooler included. Compensating the chromaticity reduces the tune spread and can help to avoid resonances. There are sextupoles foreseen in ELENA for this, but during the commissioning they were not used. We believe it to be worth trying.

[1] G. E. Lee-Whiting, End effects in first-order theory of quadrupole lenses, Nucl. Instrum. Methods **76**, 305 (1969).

[2] M. Berz, B. Erdélyi, and K. Makino, Fringe field effects in small rings of large acceptance, Phys. Rev. Accel. Beams **3**, 124001 (2000).

[3] E. Forest and J. Milutinovic, Leading order hard edge fringe fields effects exact in $(1 + \delta)$ and consistent with maxwell's equations for rectilinear magnets, Nucl. Instrum. Methods Phys. Res., Sect. A **269**, 474 (1988).

[4] B. D. Muratori, J. K. Jones, and A. Wolski, Analytical expressions for fringe fields in multipole magnets, Phys. Rev. Accel. Beams **18**, 064001 (2015).

[5] L. Bojtár, Efficient evaluation of arbitrary static electromagnetic fields with applications for symplectic particle tracking, Nucl. Instrum. Methods Phys. Res., Sect. A **948**, 162841 (2019).

[6] V. Chohan *et al.*, *Extra Low ENergy Antiproton (ELENA) ring and its Transfer Lines: Design Report* (CERN, Geneva, 2014).

[7] J. U. Brackbill and D. C. Barnes, The effect of nonzero $\nabla \cdot \mathbf{B}$ on the numerical solution of the magnetohydrodynamic equations, J. Comput. Phys. **35**, 426 (1980).

[8] C. E. Mitchell and A. J. Dragt, Accurate transfer maps for realistic beam-line elements: Straight elements, Phys. Rev. Accel. Beams **13**, 064001 (2010).

[9] D. T. Abell, Numerical computation of high-order transfer maps for rf cavities, Phys. Rev. Accel. Beams **9**, 052001 (2006).

[10] M. Venturini and A. J. Dragt, Accurate computation of transfer maps from magnetic field data, Nucl. Instrum. Methods Phys. Res., Sect. A **427**, 387 (1999).

[11] P. Chang, A general method for symplectic particle tracking in a three-dimensional magnetic field, Tech. Report No. SLAC-PUB-8514, 2000.

[12] J. Bahrdt, M. Scheer, and G. Wustefeld, Tracking simulations and dynamic multipole shimming for helical undulators, in *Proceedings of the 10th European Particle Accelerator Conference, Edinburgh, Scotland, 2006* (EPS-AG, Edinburgh, Scotland, 2006).

[13] M. Titze, J. Bahrdt, and G. Wüstefeld, Symplectic tracking through straight three dimensional fields by a method of generating functions, Phys. Rev. Accel. Beams **19**, 014001 (2016).

[14] F. Mackay, R. Marchand, and K. Kabin, Divergence-free magnetic field interpolation and charged particle trajectory integration, J. Geophys. Res.: Space Phys. **111** (2006).

[15] A. J. Dragt, *Lie Methods for Nonlinear Dynamics with Applications to Accelerator Physics* (University of Maryland Press, College Park, 2018).

[16] P. Belochitskii and L. V. Joergensen (private communication).

[17] M. Tao, Explicit symplectic approximation of nonseparable hamiltonians: Algorithm and long time performance, Phys. Rev. E **94**, 043303 (2016).

[18] H. Grote and F. Schmidt, Mad-x users guide (2019), http://madx.web.cern.ch/madx/.

[19] K. Atkinson and W. Han, *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction* (Springer, Berlin, Heidelberg, 2012).

[20] E. Forest, F. Schmidt, and E. McIntosh, Introduction to the polymorphic tracking code, KEK Report No. 2002-3.

[21] F. Asvesta, N. Karastathis, and P. Zisopoulos, PyNAFF: A Python module that implements NAFF, https://github.com/nkarast/PyNAFF.

[22] K. L. Brown, *A First- and Second-order Matrix Theory for the Design of Beam Transport Systems and Charged Particle Spectrometers* (SLAC, Stanford, CA, 1972).

[23] W. Scandale, Dynamic aperture, in *CERN Accelerator School: Course on Advanced Accelerator Physics (CAS)* (CERN, Geneva, 1995), pp. 0109–146.

[24] J. Laskar, Secular evolution of the solar system over 10 million years, Astron. Astrophys. **198**, 341 (1988).

[25] R. Bartolini, A. Bazzani, M. Giovannozzi, W. Scandale, and E. Todesco, Tune evaluation in simulations and experiment, Part. Accel. **52**, 147 (1996).

[26] R. Bartolini, M. Giovannozzi, W. Scandale, A. Bazzani, and E. Todesco, Precise measurement of the betatron tune, Part. Accel. **55**, 1 (1996).

[27] D. Robin, C. Steier, J. Laskar, and L. Nadolski, Global Dynamics of the Advanced Light Source Revealed through Experimental Frequency Map Analysis, Phys. Rev. Lett. **85**, 558 (2000).

[28] J. Laskar, Frequency analysis for multi-dimensional systems: Global dynamics and diffusion, Phys. Rev. D **67**, 257281 (1993).

[29] J. Laskar, The chaotic motion of the solar system: A numerical estimate of the size of the chaotic zones, Icarus **88**, 266 (1990).

[30] E. Todesco, M. Giovannozzi, and W. Scandale, Fast indicators of long term stability, Part. Accel. **55**, 27 (1996).

[31] E. Todesco and M. Giovannozzi, Dynamic aperture estimates and phase-space distortions in nonlinear betatron motion, Phys. Rev. E **53**, 4067 (1996).

[32] M. Giovannozzi and E. Todesco, Numerical methods to estimate the dynamic aperture, Part. Accel. **54**, 203 (1996).

[33] J. D. Meiss, Symplectic maps, variational principles, and transport, Rev. Mod. Phys. **64**, 795 (1992).

[34] M. Giovannozzi, W. Scandale, and E. Todesco, Dynamic aperture extrapolation in presence of tune modulation, Phys. Rev. E **57**, 3432 (1998).

[35] M. Giovannozzi, W. Scandale, and E. Todesco, Inverse logarithmic extrapolation of survival plots in hadron colliders, in *Beam Dynamics Newsletters*, Vol. 12 (1996), p. 6.

[36] M. Giovannozzi, W. Scandale, and E. Todesco, Prediction of long term stability in large hadron colliders, Part. Accel. **56**, 195 (1997).