Use of transfer maps for modeling beam dynamics in a nonscaling fixed-field alternating-gradient accelerator

Y. Giboudot

University of Brunel, Uxbridge, United Kingdom

A. Wolski*

University of Liverpool, Liverpool, United Kingdom

(Received 8 October 2011; published 3 April 2012)

Transfer maps for magnetic components are fundamental to studies of beam dynamics in accelerators. In the work presented here, transfer maps are computed in Taylor form for a particle moving through any specified magnetostatic field by applying an explicit symplectic integrator in a differential algebra code. The techniques developed are illustrated by their application to study the beam dynamics in the electron model for many applications (EMMA), the first nonscaling fixed-field alternating-gradient accelerator ever built. The EMMA lattice has 4 degrees of freedom (strength and transverse position of each of the two quadrupoles in each periodic cell). Transfer maps may be used to predict efficiently the dynamics in any lattice configuration. The transfer map is represented by a mixed variable generating function, obtained by interpolation between the maps for a set of reference configurations: use of mixed variable generating functions ensures the symplecticity of the map. An optimization routine uses the interpolation technique to look for a lattice defined by four constraints on the time of flight at different beam energies. This provides a way to determine the lattice configuration required to produce the desired dynamical characteristics. These tools are benchmarked against data from the recent EMMA commissioning.

DOI: [10.1103/PhysRevSTAB.15.044001](https://doi.org/10.1103/PhysRevSTAB.15.044001)

PACS numbers: 29.20.D-, 29.27.Bd

I. INTRODUCTION

Modeling the dynamics of particles in accelerators is essential for optimizing design and performance. Determining the trajectories of particles through the magnetic fields in an accelerator is fundamental to the dynamics; a wide variety of phenomena, including beam lifetime, spin polarization, and synchrotron radiation emission depend directly on the particle trajectories. To aid the modeling process, transfer maps for standard accelerator components, such as solenoids, dipoles, quadrupoles, and higher-order multipoles, have been developed and used for almost as long as accelerators themselves (see, for example, [1]).

However, expressing a map in closed form for even some relatively simple accelerator components requires a number of assumptions and approximations. Fringe fields and high-order multipoles, when not ignored altogether, are often represented by simplified models. But as accelerators begin to operate in more demanding regimes of beam quality and stability, it becomes important to provide descriptions of particle trajectories through

complex magnetic field configurations that are as complete and accurate as possible.

The development of the techniques described in this paper was motivated by the need for efficient, accurate modeling of beam dynamics in EMMA [2–5], a nonscaling fixed-field alternating-gradient accelerator (FFAG). Such machines are of interest for applications such as muon acceleration [6], and hadron cancer therapy [7–9], and have been the subject of numerous design studies and beam dynamics simulations [10–14].

In EMMA, a highly compact doublet cell is achieved using short quadrupole magnets [15]. A large aperture requirement then leads to fringe fields that potentially have a significant impact on the beam dynamics. Accurate simulations of the beam dynamics in EMMA require a dense description of the magnetic field over a large volume, and numerous integration steps. Solving Maxwell's equations in an EMMA cell (by a finite element code, OPERA [16]) leads to a 3D magnetic field map that can be used for tracking in EMMA with a numerical code such as PYZGOUBI [17–20], that integrates the equations of motion in small steps through the field. In most cases, PYZGOUBI routines are fast and reliable. However, in a purely numerical approach, the insight into the dynamics provided by an analytical map is lost. Furthermore, the field data can be cumbersome large, and numerical integration is often too slow for studies where large numbers of particle turns are required, such as for lattice optimization studies.

*a.wolski@liverpool.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 3.0 License. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

An alternative approach is to construct a transfer map in some standard form, such as a Taylor series, with coefficients derived from numerical field data. Such a map, if including terms to sufficiently high order, can capture the important features of the dynamics, while providing a more compact and efficient representation of an accelerator component, or set of accelerator components. Furthermore, it is often possible, by inspecting coefficients in a Taylor map describing the dynamics, to determine without the need for further tracking a variety of dynamical properties, such as betatron tunes, tune shifts with amplitude, chromaticity, and coupling.

However, representing the dynamics by a finite-order Taylor map raises some well-known issues. In particular, truncating the Taylor map at some given order risks limiting the accuracy of the map. Although the accuracy of the map may be improved, in principle, by including higher-order terms, it can be difficult to compute terms of very high order with good accuracy. Also, truncation of a Taylor map can lead to loss of symplecticity, which can be important in some situations (for example, when modeling the dynamic aperture of a storage ring, or attempting to characterize some weak nonsymplectic effects). In an FFAG such as EMMA, the beam will generally be accelerated over a small number of turns, and loss of symplecticity is therefore not likely to be a critical issue. However, the loss of accuracy associated with the truncation of a Taylor map can be important, since the energy and trajectory of each particle will cover a large region in phase space during the acceleration process. A possible solution is to compute a set of transfer maps, covering the range of energies and trajectories traversed by a particle during acceleration: the values of the dynamical variables can then be kept within some specified limits. The question then arises as to the number of transfer maps required (or, equivalently, the range over which the accuracy of each map can be considered acceptable).

In this paper, we report the application of methods intended to provide fast and accurate tools for the analysis of beam dynamics in EMMA. The basis of our technique is the construction of high-order transfer maps in the form of Taylor series and mixed variable generating functions, using detailed magnetic field data. Transfer maps have been used for the analysis and design of accelerator beam lines for many years; see, for example, [1,21–29]. We build, in particular, on the work of Dragt and his collaborators, in constructing transfer maps from detailed magnetic field data [30–32]. Our tools for constructing arbitrary-order Taylor maps for the particle dynamics are based on the use of a differential algebra code, COSY [27,33–36], to integrate the equations of motion for a particle moving in a magnetic field; the field is represented, using Dragt’s techniques, as a set of modes fitted to a detailed magnetic field map. Our first goal is to investigate the accuracy of the transfer map when applied to EMMA.

As a nonscaling FFAG, EMMA provides a novel regime for application of these techniques: the particle trajectories cover a large range during acceleration (with the beam energy increasing by a factor of 2); the fringe field effects of the magnets are important; and the dynamics must be described accurately in the longitudinal as well as the transverse degrees of freedom.

Our second goal is to develop our modeling tools to describe the effects of varying the machine configuration. EMMA has the unusual feature that the strengths and horizontal positions of each of the two quadrupoles forming a periodic cell can be changed, to provide control over the variation of parameters such as the betatron tunes and time of flight with energy. Characterizing each new lattice configuration using numerical integration is a slow process: we show how transfer maps, once calculated for a set of reference lattice configurations, can be applied to determine accurately and efficiently the properties of any specified new configuration. This allows a model to be fitted to machine data, and we show that the fitted model can be used to make accurate predictions of the effect of configuration changes on the time of flight. A key step in this process is the novel technique of constructing a map by interpolation between maps (in the form of mixed variable generating functions) representing the dynamics in a specified set of reference lattice configurations.

The paper is organized as follows. In Sec. II, we describe the methods used, including all the main steps from design of the magnets to computation of the beam dynamics. In Sec. III, we present an application of the method to explore some dynamical properties of the nonscaling FFAG, EMMA. Comparisons are given, for transverse and longitudinal dynamics, between the results of a numerical tracking code (PYZGOUBI) and calculations based on a transfer map. One of the main goals of this section is to understand the limits on accuracy of the transfer map; although the results strictly apply to the case of a single machine (EMMA), the challenges of this particular case are such that the results indicate the potential use of the technique for many other machines. In Sec. IV we describe how we have used transfer map techniques to make possible an optimization tool that can be used for finding the lattice configuration that most closely achieves specified dynamical properties. Finally, in Sec. V, we compare the simulation results with experimental results from the early stages of commissioning of EMMA.

II. CONSTRUCTION OF TRANSFER MAPS

Construction of a transfer map for a given accelerator section is performed in three stages. First, the design of the section is used in a magnetic field solver, to construct a detailed 3D magnetic field map of the section. This field map is in numerical form, that is, it consists of values for the three field components over a grid of spatial coordinates. The second stage is to convert this numerical field

map into a semianalytical form, in which the field is described as a superposition of a set of modes, each mode expressing the field as a function of the coordinates. The mode coefficients provide the input for the third stage, in which a differential algebra code is used to integrate the equations of motion of a particle moving through the field, with the dynamical variables expressed at each stage of the integration as functions (Taylor series) of the initial values of the dynamical variables. The result of the integration is a transfer map for the specified section. In this section, we describe each of these stages in more detail, using the computation of a transfer map for a single EMMA cell for illustration.

A. Magnetic field computation

Each cell within the EMMA ring includes a horizontally focusing (F) quadrupole magnet, and a horizontally defocusing (D) quadrupole magnet (see Fig. 1). EMMA consists of 42 cells; therefore, the periodicity condition involves an exit face rotation by an angle $2\pi/42$ to match the following entrance face (which is parallel to the faces of the magnets in the following cell). All the elements are

FIG. 1. Schematic top view of an EMMA cell. The magnetic axes of the D and F magnets are offset horizontally (by distances x_d and x_f , respectively) with respect to the 42-sided reference polygon used to define the machine layout.

aligned with respect to a straight line of length 394.481 mm, which represents one side of a 42-sided polygon. This polygon is defined so as to be close to the expected closed orbit of a particle with energy 15 MeV (the machine is designed to accelerate electrons from 10 to 20 MeV). The vacuum chamber extends transversely a few centimeters on each side: simulations of the magnetic field and the particle motion have to be precise in this area. A 3D model of the magnets, required to provide an accurate description of the magnetic field, was developed in OPERA [16], which uses the finite element method to solve Maxwell's equations for given current and magnetic material distributions.

The magnetic lattice in EMMA has 4 degrees of freedom, corresponding to the strengths and horizontal positions of each of the two quadrupoles within the cell. Because each magnet in a cell is within the fringe field region of the other, changing the strength or position of one magnet can affect the field produced by the other. As long as the iron in the yokes of the magnets is not close to saturation (so that the field varies linearly with the current), simulations suggest that it is possible to construct the field map for any given quadrupole strengths by scaling and superposition of the field maps for each quadrupole powered separately, but taking into account the presence of the yoke of the other magnet [20,37]. Then, only field maps for the quadrupoles with a range of different relative positions (and not for a range of different currents) need to be computed in OPERA.

B. Analytical field representation

The numerical field map obtained using OPERA is suitable for tracking in a numerical code, such as PYZGOUBI. However, our goal is to construct a transfer map that represents the dynamics for a complete cell: for this, as we shall see, an analytical representation of the field is needed. In this section, we shall describe how we construct an analytical representation of the field from the numerical field map. Briefly, we write down a set of modes for the field (capturing the shape in 3D), and express the field as a set of coefficients for the different modes. This process has a number of advantages, beyond the provision of a representation allowing the construction of a transfer map. First, the original field data consist of the field components given (typically) on a 3D Cartesian grid with 1 mm step in all directions: the data files for field maps in this form tend to be rather large and cumbersome. A set of mode coefficients, on the other hand, provides a much more compact and convenient representation of the field. Second, detailed inspection of the data to identify errors can be difficult. However, in a mode decomposition, Maxwell's equations are enforced for any values of the coefficients, by the form of the basis functions. Errors in the field map can often be identified simply by inspection of the coefficients since, for an appropriate choice of basis, the various coefficients can

be identified with symmetries in the system generating the magnetic field.

In the present case, we are interested in constructing the transfer map for a particle traveling within the vacuum chamber passing through one complete cell. The coefficients of a mode decomposition of a magnetic field can be accurately determined by fitting the field on a closed surface within the field. Within the surface, residuals to the fit decrease exponentially; outside the surface, they increase exponentially. Therefore, it is beneficial to choose a surface that encloses the greatest possible volume. Ideally, in the case of EMMA, we would use the surface defined by the inner surface of the vacuum chamber for performing the fit; however, this has a rather complicated geometry. To simplify the fit, we use a cylinder with the maximum radius such that it just fits within the vacuum chamber (which itself just fits within the apertures of the magnets). Although a larger volume could be enclosed by choosing a volume with elliptical cross section, the basis functions for an elliptical cross section involve more complicated functions [32]; since a circular cross section provides an adequate region for our purposes, we use a cylinder with circular cross section, with the axis of the cylinder collinear with the reference polygon used to define the layout of the machine (see Fig. 1). An appropriate set of basis functions for the magnetic field is expressed in cylindrical polar coordinates, as follows:

$$B_\rho = \sum_{m,n} a_{mn} I'_m(nk_z \rho) e^{im\phi} e^{ink_z z}, \quad (1a)$$

$$B_\phi = i \sum_{m,n} a_{mn} \frac{m}{nk_z \rho} I_m(nk_z \rho) e^{im\phi} e^{ink_z z}, \quad (1b)$$

$$B_z = i \sum_{m,n} a_{mn} I_m(nk_z \rho) e^{im\phi} e^{ink_z z}. \quad (1c)$$

Here, $I_m(nk_z \rho)$ are modified Bessel functions of the second kind. Note that the fields represented by these functions satisfy Maxwell's equations, and satisfy also the rotational symmetry $\vec{B}(\phi + 2\pi) = \vec{B}(\phi)$, as is required for any physical fields.

The coefficients a_{mn} are obtained by performing a 2D discrete Fourier transform on the values of B_ρ on a cylinder of radius ρ_0 , whose axis is coincident with the z axis. A given a_{mn} is simply the corresponding Fourier coefficient, normalized by $I'_m(nk_z \rho_0)$. Longitudinally, the analysis is simplified if the model edges (at $z = 0$ and $z = 394.481$ mm) are far enough from the magnets that the fringe fields vanish on both sides; in this case, a reduced set of basis functions can be used. As we will see, in reality, the field does not go lower than a few gauss between two cells. As mentioned above, the residuals of the fit increase exponentially with increasing radius ρ : if the fit achieves a certain level of accuracy on the surface of the cylinder $\rho = \rho_0$, then the accuracy will be better than this for $\rho < \rho_0$ (i.e. within the cylinder), but can become rapidly worse for $\rho > \rho_0$. In practice, this means that when tracking particles

through the field represented by a mode expansion, care should be taken to avoid initial conditions that lead to trajectories crossing the surface on which the field fit was performed.

Figure 2 shows the radial field component and the residual of the fit on a reference cylinder of radius $\rho_0 = 15$ mm in one EMMA cell. The relatively large residuals at the entrance and exit of the cell arise from the nonzero values of the field, which cannot be represented by the reduced set of Fourier basis functions we have used. It would be possible to include these fields using the complete set of basis functions, at the cost of an increased number of coefficients, and a corresponding increase in computation time for the integration of the equations of

FIG. 2. (a) Radial component of the magnetic field and (b) residual of the fit on a reference cylinder of radius 15 mm within one EMMA cell.

motion; however, with errors of the order of few gauss, this fit is considered good enough for our tracking studies.

Although cylindrical polar coordinates are appropriate for the field description, beam dynamics studies are more conveniently performed using Cartesian coordinates. Therefore, we obtain a representation of the field using basis functions in Cartesian coordinates: this is the basis used for the construction of the transfer map, described in Sec. II C below. In the Cartesian basis, the components of the magnetic field are written

$$B_x = -i \sum_{m,n} c_{mn} \frac{mk_x}{k_y} e^{imk_x x} e^{-k_y y} e^{ink_z z}, \quad (2a)$$

$$B_y = \sum_{m,n} c_{mn} e^{imk_x x} e^{-k_y y} e^{ink_z z}, \quad (2b)$$

$$B_z = -i \sum_{m,n} c_{mn} \frac{nk_z}{k_y} e^{imk_x x} e^{-k_y y} e^{ink_z z}. \quad (2c)$$

If the parameters k_x , k_y , and k_z are related by

$$k_y^2 = m^2 k_x^2 + n^2 k_z^2, \quad (3)$$

then the fields given in Eq. (2) satisfy Maxwell's equations. In principle, a fit of the Cartesian basis to the numerical field data can be obtained by performing a discrete Fourier transform of the vertical field component B_y on a plane $y = y_0$. Note that the fit has the property that the residuals increase exponentially with increasing y , corresponding to the exponential increase with increasing ρ in the cylindrical case. It is therefore desirable to choose y_0 as large as possible: however, for magnets with a circular aperture, the width of the plane (i.e. the maximum value of x) is reduced with increasing y_0 , and the maximum volume enclosed by the surface of the fit is therefore smaller than if a cylindrical geometry is used. In the case of EMMA, the situation is a little more complicated, since the two quadrupoles in a cell are aligned with a horizontal offset between their axes: it is therefore difficult to fill the vacuum chamber even using a cylindrical geometry. However, a fit based on a rectangular geometry has a further disadvantage, in that the value of the horizontal mode number [k_x in Eqs. (2) and (3)] is not determined by the geometry. This is in contrast to the cylindrical basis, where rotational periodicity requires that the azimuthal mode numbers m in Eq. (1) be integers. The Cartesian basis in Eq. (2) implicitly assumes a periodicity in x , which does not really exist. Essentially, the Cartesian basis is less appropriate to the geometry of an accelerator beam line, and generally, a fit to numerical field data based on Eq. (2) is less successful than one using cylindrical polar coordinates, Eq. (1). However, in practice it is found to be possible to obtain a reasonable description of the field by converting the coefficients a_{mn} (describing the cylindrical modes) to an equivalent set c_{mn} (describing the Cartesian modes). The formulas for performing the conversion are presented and derived in the Appendix.

The value for k_x in the Cartesian basis is not determined *a priori*; however, it can be chosen to minimize the residual in the final field description, compared to the original numerical data. Figure 3 shows such a comparison, with the coefficients in the Cartesian expansion obtained using Eq. (A7) from the coefficients in the cylindrical expansion (illustrated in Fig. 2). Mode numbers up to 10 in the transverse direction and up to 40 in the longitudinal direction have been included. The residuals can be minimized using a value of k_x anywhere in the range $10 \text{ m}^{-1} < k_x < 30 \text{ m}^{-1}$; outside this range, the residuals start to increase, but the dependence on the value of k_x is not very sensitive. A value of 30 m^{-1} was used for the plots in Fig. 3. It can be seen in Fig. 3(b) that there is a tendency for the residuals to

FIG. 3. Residuals of the fit to the field in a single EMMA cell (a) in the median plane, and (b) in the plane with coordinate $y = 10 \text{ mm}$. In both cases, the residuals increase with the horizontal excursion x (i.e. approaching the surface on which the fit is performed).

increase for x approaching ± 10 mm: at $y = 10$ mm, points with $|x| > 11$ mm are outside the surface on which the fit to the field was performed. Note that the residuals have a structure such that they appear to oscillate as a function of longitudinal position. The effect of the residual integrated over the length of the cell is small compared to the overall effect of the fields of the magnets. It should be possible to reduce the residuals still further by including higher-order longitudinal modes; however, the residuals are already small enough that they do not have a significant impact on the dynamics. On the midplane, the residual reaches a maximum of 2 G at the entrance of the D quadrupole. Since the field at that point is approximately 2000 G, the relative error is 0.1%. For $y = 10$ mm, the residual has a maximum of 4 G which is also small compared to the maximum field strength.

C. Construction of a transfer map

Numerically, the particle dynamics can be obtained by integrating the equations of motion for a particle in the field, using an appropriate integration algorithm. For example, a Runge-Kutta integrator could be used. However, since symplecticity is a desirable property, we use an explicit symplectic integrator. A convenient algorithm has been developed using Lie algebra techniques by Wu, Forest, and Robin (“WFR integrator” [38]): this is applicable to the case of a particle moving through a magnetic field where the field varies with longitudinal position, and is the algorithm we use for the studies of beam dynamics in EMMA.

A numerical integrator “solves” the dynamics in the sense that if one is given the initial values of the dynamical variables for a particle at the entrance to a cell, then by applying the integrator one can compute the final values (the values of the dynamical variables for the particle at the exit of the cell). This only requires knowledge of the magnetic field at particular locations along the trajectory of the particle. However, the result gives no indication of how changes in the initial values will affect the final values: the integrator must be run separately for each set of initial values that may be of interest, and this can be very time consuming. If an investigation requires tracking many particles many times (e.g. if investigating the effect of configuration changes in the magnetic lattice), then a more efficient technique may be needed. One solution is to construct a transfer map, which expresses in functional form the dependence of the final values of the dynamical variables, in terms of the initial values. A transfer map can be constructed by implementing the same algorithm used for numerical integration, but using differential algebra (DA) variables, rather than numerical variables. Each DA variable contains a power series up to some order, in terms of specified variables (in our case, the dynamical variables). Since each step of the integration essentially propagates a range of starting conditions, the integrator does

not trace the path of a particle along a single trajectory; therefore, to carry out the integration, the dependence of the magnetic field on the spatial coordinates must be expressed in functional form. This can be achieved using a mode decomposition, as described in the previous section. The result of the integration is a power series, or rather a set of power series (one for each dynamical variable). Once the transfer map has been obtained (by a single integration through the cell), tracking particles simply involves evaluating the power series for the given initial values of the dynamical variables. This process is much more efficient than repeatedly integrating the equations of motion every time different initial conditions are used.

We have implemented the WFR integrator in the DA code COSY [36] to construct the transfer map for a particle moving through a single cell in EMMA, for any given configuration of the lattice. The dynamical variables are as follows: x and y are the coordinates in a plane transverse to the reference trajectory at any point along the reference trajectory. p_x and p_y are the normalized momenta canonically conjugate to x and y , given by

$$p_x = \frac{\gamma m \dot{x} + q A_x}{P_0}, \quad (4)$$

$$p_y = \frac{\gamma m \dot{y} + q A_y}{P_0}, \quad (5)$$

where γ is the relativistic factor for the particle, m and q are the mass and charge of the particle, A_x and A_y are the transverse components of the electromagnetic vector potential, and P_0 is a fixed reference momentum (related to the reference energy E_0 by $E_0^2 = P_0^2 c^2 + m^2 c^4$). A dot indicates the derivative with respect to time. The longitudinal coordinate z of a particle is defined by

$$z = s/\beta_0 - ct, \quad (6)$$

where β_0 is the normalized velocity of a particle with the reference momentum, and t is the time at which the particle arrives at a point s along the reference trajectory. The energy deviation δ , canonically conjugate to z , is defined by

$$\delta = \frac{E}{P_0 c} - \frac{1}{\beta_0}, \quad (7)$$

where E is the energy of the particle. With these variables, the Hamiltonian describing the dynamics of a charged particle is

$$H = \frac{\delta}{\beta_0} - \sqrt{\left(\frac{1}{\beta_0} + \delta\right)^2 - \left(p_x - \frac{q A_x}{P_0}\right)^2 - \left(p_y - \frac{q A_y}{P_0}\right)^2} - \frac{1}{\beta_0^2 \gamma_0^2} - \frac{q A_z}{P_0}, \quad (8)$$

where A_z is the component of the vector potential in the direction of the (straight) reference trajectory, and $\gamma_0 = 1/\sqrt{1 - \beta_0^2}$.

The magnetic field (actually, the magnetic vector potential) is given as a mode expansion in a Cartesian basis, obtained using the methods described in the previous section. Since the WFR integrator requires the paraxial approximation, some care is needed in its application, since in an FFAG (for example), beam excursions can be large (of the order of a few cm), and the paraxial approximation may not be valid. In principle, it is possible to make some estimate of the closed orbit at a given reference energy (for example, by using a simple linear model, with hard-edge magnets, for the lattice), and then to use the closed orbit as the reference trajectory. However, the curvature of the reference trajectory would then add considerable complexity to the integration of the equations of motion. Instead, we can use for the reference trajectory a straight line chosen to minimize the deviation of the (estimated) closed orbit from the reference trajectory. At least for small energy deviations, particle trajectories should then remain reasonably close to the reference trajectory. Some iteration may be needed to find the reference trajectory that minimizes the deviation along the closed orbit, when a realistic field model is used. Since each cell in EMMA bends the trajectory by $2\pi/42$, using a straight reference trajectory means that angles of order $2\pi/42$ in the closed orbit may be expected; it is not immediately clear whether the paraxial approximation will be valid under these conditions. One way to check, is to compare the results with a computation that does not make the paraxial approximation: for example, using PYZGOUBI. Since we find reasonable agreement in dynamical quantities (in particular, closed orbits, tunes, and times of flight) between the transfer map calculations and PYZGOUBI (using the magnetic field map), we conclude that it is acceptable to use the paraxial approximation, although this may be the cause of some of the observed discrepancies between the results using different methods.

Two important issues arise when applying the transfer maps constructed using this technique to an FFAG such as EMMA. First, although the integrator we have used is symplectic in the strict mathematical sense (in that the changes to the dynamical variables given at each step of the integration may be represented by a symplectic map), the fact that the DA code truncates the power series for each dynamical variable at a specified order at each step means that the final map is no longer symplectic. However, the map is symplectic up to the order of the power series in the sense that symplecticity may be restored by adding appropriate higher-order terms. Decreasing the step size in the integration will improve the accuracy of the final map, but will not reduce the symplectic error (conversely, using a large step size in the integration will reduce the computation time, at the cost of some loss of accuracy, but without introducing any additional symplectic error). The symplectic error can only be reduced by increasing the order of the power series manipulated by the DA code: however, increasing the order

can significantly increase the computation time. Therefore, the question arises as to what order of truncation is necessary, to reduce the symplectic error to an acceptable level. Of course, a high-order truncation may also be needed simply to achieve reasonable accuracy. The question as to the appropriate order of truncation will be explored, for the case of EMMA, in Sec. III.

A second issue concerns the number of maps required to give an accurate description of the dynamics over the full parameter range. As already mentioned, to ensure the validity of the paraxial approximation for EMMA, we can use a number of separate transfer maps covering the full energy range, from 10 to 20 MeV. Each map is computed for a different reference energy and reference trajectory. Using a large number of maps should improve accuracy, by reducing the maximum values of the dynamical variables for which the map needs to be evaluated. However, a large number of maps also increases the overall computation time and reduces the convenience of the technique. This issue (which is related to the question of the appropriate truncation order) is also discussed in Sec. III.

III. BEAM DYNAMICS IN EMMA USING TRANSFER MAPS

A. Comparison between transfer map and numerical tracking

As an initial test of the transfer maps computed using the techniques described in the previous section, we computed some of the key features of the beam dynamics in EMMA, and compared the results with those obtained using PYZGOUBI (a numerical tracking code). In particular, we looked at the closed orbit, horizontal and vertical tunes, and time of flight. All quantities were computed assuming fixed beam energy (i.e. no acceleration), but were computed as a function of energy. One of the interesting aspects of the EMMA design is the impact of fringe fields on the beam dynamics. We therefore used PYZGOUBI to compute the features of the dynamics not only with a magnetic field map for the quadrupoles including the fringe fields, but also with a “hard-edge” model for the magnets, in which the field was represented as falling abruptly to zero at the entrance and exit faces of the magnets.

The closed orbit is found simply by searching for the trajectory (at fixed energy) for which the phase space variables return, at the exit of a single cell, to their values at the entrance of the cell. For numerical integration (using PYZGOUBI), tracking a particle starting on the reference trajectory gives the zeroth order term (represented by a phase space vector $\vec{x}_0$) in the map; and tracking particles with small deviations from the reference trajectory allows the construction of the first-order terms in the map (represented by a matrix M). If $\vec{x}_i$ is the initial phase space vector, the final phase space vector $\vec{x}_f$ is then given in a linear approximation by

$$\vec{x}_f = M\vec{x}_i + \vec{x}_0. \quad (9)$$

The closed orbit is defined by $\vec{x}_{co} = M\vec{x}_{co} + \vec{x}_0$, and, hence,

$$\vec{x}_{co} = (I - M)^{-1}\vec{x}_0, \quad (10)$$

where I is the identity matrix. The presence of nonlinear terms in the map generally means that some iteration is required to find the closed orbit with good accuracy. In the case of a transfer map, if the transfer map is constructed around a given reference trajectory, then the closed orbit at the entrance to a cell can be found by searching for the fixed point of the map. Integrating through a cell in a differential algebra code (COSY) produces the zeroth and first-order terms in the map directly, in which case the closed orbit can be estimated by the above formulas. Again, some iteration is needed to identify the closed orbit with good accuracy: the iteration is performed using a single transfer map, though some adjustment of the reference trajectory may be appropriate to minimize the closed orbit excursion (in which case a new map is computed around the new reference trajectory). Some care must be taken with the coordinate rotation at the exit of the cell. Standard formulas exist for the transformation associated with such a rotation [26]: however, these formulas apply to rotation in a field-free region. In the case of EMMA, because of the fringe field of the magnets, it is necessary when constructing the closed orbit to define a cell such that the exit of the cell is some distance from the face of a magnet. Therefore, for the reference trajectory, we use a modification of the 42-sided polygon, in which each side is extended back from the vertex, to the center of the drift between the magnets. This means that the transformation at the exit of each cell involves a (horizontal) translation and a drift, as well as a rotation.

In general, we find good agreement for the closed orbit from PYZGOUBI (using a magnetic field map) and from COSY: differences are typically of order $10 \mu\text{m}$. Obtaining the closed orbit accurately is important, since it is the variation of the closed orbit with energy that dominates the variation of the time of flight with energy; and the variation of the time of flight in turn is important for achieving acceleration. Figure 4 shows the closed orbit at selected reference energies from 10 to 20 MeV, found using the transfer map for a particular lattice configuration. Note that the increasing average radius with energy is compensated by a “smoothing out” of the closed orbit: this leads to a parabolic variation of time of flight with energy, with a minimum time of flight around 14 MeV.

After finding the closed orbit, the next step is to study the behavior of particles around this closed orbit. Betatron motion was studied by applying the transfer maps to particles with some initial transverse offset with respect to the closed orbit. Figure 5 (top) shows the horizontal phase space for reference energies from 10 to 20 MeV, constructed by applying the corresponding transfer map

FIG. 4. Closed orbit in a single EMMA cell, at selected reference energies from 10 to 20 MeV. The reference trajectory [$x(s) = 0$] is a modification of the 42-sided polygon, in which each side is extended back from a vertex to the middle of the drift between the magnets. The complete transfer map for a cell then involves a coordinate transformation at the exit of the cell, consisting of a horizontal translation and a rotation: the effect of this transformation is to restore the position and angle of each trajectory at the exit of the cell (the right-hand edge of the plot) to the position and angle at the entrance of the cell (the left-hand edge of the plot).

(in energy) for 2000 iterations to particles with 1 mm initial transverse offset. We notice that there is some nonphysical growth in the amplitude over time for 10 and 12 MeV: this is a consequence of a symplectic error introduced by truncation of the transfer map to 2nd order. The effect becomes significantly less visible if terms up to 3rd order are retained: see Fig. 5(b).

The betatron tunes can be obtained directly from the transfer map, by computing the eigenvalues λ_k of the linear part of the map. The tunes ν_k are given by

$$\nu_k = \frac{|\ln \lambda_k|}{2\pi}. \quad (11)$$

The betatron tunes can be obtained in a similar way from the numerical tracking: the linear part of the map is first obtained by tracking particles with small deviations (in each of the transverse dynamical variables), and then the tunes are found from the eigenvalues of the linear part of the map. Figure 6 shows the variation of horizontal tune, vertical tune, and time of flight with energy, computed using each of the various techniques and models (transfer map, numerical tracking through a field map, and numerical tracking in a hard-edge magnet model). Looking first at the comparison between the numerical tracking through a field map (red crosses) and the transfer map (blue circles), we observe a discrepancy smaller than the expected

FIG. 5. Betatron motion at different reference energies with transfer maps (a) of order 2, and (b) of order 3. A different map was computed for each reference energy. For the maps of order 2, there is some nonphysical growth in the amplitude over time for 10 and 12 MeV: this is a consequence of the loss of symplecticity arising from truncation of the map. The effect becomes significantly less visible if terms up to 3rd order are retained.

measurement precision for the vertical tune and path length. However, we find a larger discrepancy for the horizontal tune at low energy: 0.04 at 10 and 0.02 at 11 MeV. An explanation for this discrepancy could lie in the large excursion of closed orbits for lower energies. At such large excursions, the analytical representation of the magnetic field is at the limit of its validity, and may therefore be slightly different from the original numerical field map used to track with PYZGOUBI.

Tracking results using a hard-edge model of the magnets in PYZGOUBI are also presented (in green) on the plots, and

FIG. 6. (a) Horizontal tune, (b) vertical tune, and (c) time of flight versus kinetic energy computed with a hard-edge model, and with field maps. We observe a maximum discrepancy of about 0.05 at 10 MeV for the horizontal tune per cell, between the transfer map and numerical tracking with the fringe field, but generally the two techniques are in good agreement. However, there are significant differences over the full energy range between the hard-edge model and the detailed fringe field (field map) model. With fringe fields, the minimum time of flight is reduced by 200 ps, and the energy at which this occurs is shifted from 15 to 17 MeV.

indicate the impact of the fringe field: we observe a maximum discrepancy of about 0.05 at 10 MeV for the horizontal and vertical tune per cell with respect to the tracking using a field map with PYZGOUBI.

In Fig. 6, we observe as well a significant change in the time-of-flight curve when using field maps. The minimum of the curve is shifted from 15 MeV with hard-edge model (in agreement with the design studies) to 17 MeV, and a decrease in the minimum time of flight of about 200 ns. PYZGOUBI and transfer map simulations agree on this result, which suggests that this large discrepancy is due to the difference between the hard-edge representation and the more realistic field map representation of the EMMA cell. The baseline lattice was specifically designed to produce a time-of-flight curve that was parabolic, and centered at 15 MeV. If the models based on the field map correctly predict the real machine behavior, then a new baseline configuration should be identified, to obtain the desired time-of-flight curve (i.e. so that the minimum of the curve occurs for a specified energy, and so that the time of flight is properly adapted to the rf frequency of the accelerating cavities). This issue will be discussed further in Sec. IV.

B. Validity of transfer maps for large energy deviation

A transfer map is generally accurate only in the vicinity of a given reference trajectory, for particles traveling with an energy close to a specified reference energy. Since in EMMA the energy and the transverse excursion cover a

large range during the course of acceleration, the range of validity of a transfer map has to be examined carefully. To investigate the range of validity of a map, we can compare properties of the beam dynamics computed using a single map with a range of values of the energy deviation, with properties computed using a set of maps, each with a different reference energy.

Figure 7 shows the horizontal tune (in blue) and the vertical tune (in red) as functions of energy, obtained from 11 different transfer maps computed for different reference energies, from 10 to 20 MeV, in steps of 1 MeV (points on line). Also shown are the horizontal and vertical tunes computed from a single transfer map at a single reference energy (14 MeV), but with different values for the energy deviation δ and for different truncation order of the map: 3 (triangles), 7 (circles), and 9 (crosses). We observe a large discrepancy between the two methods for low and high energy (i.e. for large energy deviation from the reference 14 MeV). The discrepancy decreases when increasing the truncation order from 3 to 7 but no significant improvements can be seen when truncating at the 9th order; this suggests that the 7th order is the optimal truncation order for the tune variation with energy deviation.

The expected precision of the tune measurement in EMMA is 0.01. The agreement between the two methods is within this range for energy deviation up to ± 2 MeV (i.e., 15%). This indicates that it may be necessary to use at least three different maps (with reference energies of, for example, 12.5, 15, and 17.5 MeV) to reproduce the transverse and longitudinal dynamics over the whole energy range in EMMA.

In order to simulate acceleration in EMMA, the transfer maps should describe accurately the variation in time of flight with energy. Figure 8 shows the time of flight of the closed orbit at different energies computed using multiple

FIG. 7. Betatron tunes computed with a set of transfer maps at different reference energies (points on line), and with a single transfer map with different energy deviations (triangles, circles, crosses). The agreement between the two methods is better than the expected measurement precision of 0.01, for energy deviation up to ± 2 MeV (i.e., 15%). This suggests that it may be necessary to use at least three different maps (e.g. at 12.5, 15, and 17.5 MeV reference energies) to reproduce the transverse dynamics accurately over the full energy range.

FIG. 8. Time of flight for various energies computed with a set of transfer maps at different reference energies (line), and with a single transfer map with different energy deviations (triangles, squares, circles, crosses for orders 3, 4, 7, and 9, respectively). We see good agreement for the maps truncated at 4th order or above.

maps (red circles), and using a single map with energy deviation, and with different truncation orders from 3 to 9. We observe a good agreement (within the measurement precision of 10 ps) for the whole energy range for a transfer map truncated at the fourth order; no significant further improvement is obtained when increasing the truncation order up to 9. We conclude that a single map truncated at the fourth order is sufficient to describe the time-of-flight variation with energy over the whole energy range.

As a final comparison, we look at the betatron motion in horizontal phase space. We have already seen the horizontal

phase space constructed using maps of different reference energies, in Fig. 5. There, we observed that to avoid rapid unphysical growth in the betatron amplitude, it was necessary to retain terms in the map up to 3rd order. Figure 9(a) shows the horizontal phase space constructed using a single transfer map, with reference energy 14 MeV, and with different beam energies obtained by varying the energy deviation. We find that, for a map truncated at the 7th order, the betatron motion in horizontal phase space has a rapid nonphysical growth in the amplitude (arising from the symplectic error in the map) even for small energy deviation (the effects are already clear for ± 2 MeV). An improved behavior, with significantly reduced amplitude growth, is obtained for energy deviation of ± 2 MeV by retaining terms in the map up to 9th order: see Fig. 9. However, the effects of the symplectic error in the transfer map are still strongly evident for energy deviation of 4 MeV, even at 9th order. Computational limits in the COSY routine that we used did not allow us to investigate this effect for higher truncation order.

C. Nonlinear dynamics

An interesting feature of transfer maps is the fact that information about the nonlinear beam dynamics can be extracted directly, without the need for numerical analysis of tracking data. We have already seen how the tunes and time of flight depend on the beam energy: this can be described using maps computed for different reference energies (with zero energy deviation), or using a map for a particular reference energy, evaluated for different values of the energy deviation. In the latter case, treating the energy deviation as a fixed parameter (rather than as a dynamical variable), we may extract the tunes and time of flight expressed as functions of the energy deviation directly from the map.

Nonlinear effects also occur in the transverse dynamics. For instance, consider the horizontal and vertical tunes. In a linear approximation, the tunes may be computed directly from the linear part of the map [see Eq. (11)]. However, in general the tunes will vary with the amplitude of the transverse motion. The dependence of the tune on the betatron amplitude (referred to as the “tune shift with amplitude”) is captured in the map by terms of higher order in the transverse dynamical variables (x, p_x) and (y, p_y).

In a scaling FFAG (see, for example [39]), the magnetic fields are designed to maintain constant values for the betatron tunes as the beam energy varies over a wide range: this leads to strong nonlinear effects in the transverse motion. In particular, there can be large tune shifts with amplitude. In a nonscaling FFAG such as EMMA, on the other hand, the magnetic lattice is constructed out of elements for which the transverse motion is essentially linear, i.e., out of quadrupoles and drift spaces. In consequence, although there can be strong nonlinear effects when longitudinal variables are involved (in particular, large chromaticity),

FIG. 9. Betatron motion for different energy deviations, simulated with a single transfer map up to (a) 7th order, and (b) 9th order. The reference energy is 14 MeV. We observe that for a map truncated at the 7th order, the betatron motion in horizontal phase space has a strong nonphysical growth in the amplitude for energy deviation ± 2 MeV. Improved behavior is obtained for energy deviation of ± 2 MeV by retaining terms up to 9th order.

the nonlinear effects in the purely transverse dynamics will generally be rather weak; in particular, the tune shifts with amplitude are expected to be small. Nevertheless, since the magnets in EMMA do not produce “ideal” quadrupole fields, it is expected that it should be possible to observe some effects, at least in simulation. In this section, we shall calculate the tune shifts with amplitude from the transfer maps, and compare with values derived from particle tracking in a numerical code (PYZGOUBI). We shall also compare with some analytical estimates, based on the description of the field in the magnets.

For obtaining the tune shifts with amplitude from the map, it is convenient to work in action-angle variables, rather than the Cartesian variables. The horizontal action is denoted J_x , and the horizontal angle ϕ_x ; these variables are related to the Cartesian variables by

$$x = \sqrt{2\beta_x J_x} \cos \phi_x, \quad (12)$$

$$p_x = -\sqrt{\frac{2J_x}{\beta_x}} (\sin \phi_x + \alpha_x \cos \phi_x). \quad (13)$$

The transfer map is expressed in terms of a Lie transformation [29,40] with generator $g(\phi_x, J_x)$, so that

$$J_x(L) = (e^{-i\tilde{g}} J_x)_{s=0}, \quad (14)$$

$$\phi_x(L) = (e^{-i\tilde{g}} \phi_x)_{s=0}. \quad (15)$$

The tune shift with amplitude is obtained by constructing the nonlinear normal form [29,41–45]. That is, we find a transformation $\mathcal{A}$ such that

$$\mathcal{A} e^{-i\tilde{g}} \mathcal{A}^{-1} = e^{-i\tilde{g}}, \quad (16)$$

where $\tilde{g}$ is a function of the action J_x only, i.e.,

$$\tilde{g} = \sum_{l=0}^{\infty} \frac{\mu_{x,l}}{l+1} J_x^{l+1}. \quad (17)$$

The tune for a given action J_x is given by

$$\nu_x = \frac{1}{2\pi} \frac{d\tilde{g}}{dJ_x}. \quad (18)$$

Methods for constructing the normalizing transformation $\mathcal{A}$ and the normal form of the generator of the Lie transformation $\tilde{g}$ are well known (see, for example [26,27,42]). We use COSY for constructing the generator of the Lie transformation corresponding to the Taylor map, and for the normal form analysis; the results are checked with a normal form analysis performed using MATHEMATICA [46]. The generator $\tilde{g}$ is computed to third order in the betatron action.

To compare the tunes (and in particular, the variation of tune with betatron action) with the numerical model, we need a method to determine the tunes from tracking data with high precision. For this, we use the method of numerical analysis of the fundamental frequencies

(NAFF [47–53]). Briefly, we collect the horizontal coordinate x_n and canonical momentum (scaled by the reference momentum) $p_{x,n}$ at the exit of each cell, for tracking through N successive cells (so that $n = 0, \dots, N$). We then construct the normalized variables:

$$\tilde{x}_n = \frac{x_n}{\sqrt{\beta_x}} = \sqrt{2J_x} \cos \phi_{x,n}, \quad (19)$$

$$\tilde{p}_{x,n} = \sqrt{\beta_x} p_{x,n} + \frac{\alpha}{\sqrt{\beta_x}} x_n = -\sqrt{2J_x} \sin \phi_{x,n}. \quad (20)$$

(In these expressions, we assume that the action J_x is constant, but the technique still works even if there are small variations in J_x .) Then, we search numerically for a value $\mu_x = \hat{\mu}_x$ that maximizes the magnitude of

$$\sum_{n=0}^N e^{i\mu_x(n/N)} (\tilde{x}_n + i\tilde{p}_{x,n}). \quad (21)$$

The tune is given (with accuracy of order $1/N^4$) by $\nu_x = \hat{\mu}_x/2\pi$.

The advantage of the transfer map approach is that determination of the tune shift with amplitude does not require multitrack tracking with a range of different betatron amplitudes: the single integration (using DA variables), followed by normal form analysis of the Lie transformations, provides the functional dependence of the tunes on the betatron action. Obtaining the same information using a numerical tracking code requires the tracking of many particles (over a range of betatron amplitudes) through many cells, and is potentially more time consuming.

Figure 10 shows the changes in tunes (relative to the tunes at zero amplitude) with respect to changes in betatron amplitude, for 19 MeV reference energy. The black lines show the variations predicted by normal form analysis applied to the transfer map; the blue lines with circles and the red points show the changes found from NAFF of tracking data produced using (respectively) the transfer map, and a numerical code (PYZGOUBI). First, we observe that the tune shift with amplitude in all cases is small (about 0.001) over a large range of variation of the action (from 0 to 200 μm). The order of magnitude of the horizontal beta function β_x in EMMA is 1 m. Therefore the maximum transverse coordinate equivalent to an action of 200 μm is $x_{\text{max}} = \sqrt{2\beta_x J_x} = 14$ mm. From the early stages of commissioning of EMMA, it was found that to achieve efficient injection, a bunch of particles must be injected on the closed orbit with a precision of a few millimeters (at the location of a maximum for β_x). Therefore, the amplitude of oscillation remains small, and any variation in the tune resulting from betatron amplitude variations is smaller than present measurement capabilities.

Regarding the horizontal motion [Fig. 10(a)], the agreement between the NAFF of the transfer map tracking (blue circles) and the normal form analysis (black line) is reasonably good: the discrepancy is less than 2×10^{-4} for

FIG. 10. Change in (a) the horizontal tune with the horizontal action J_x , and (b) the vertical tune with vertical action J_y , for reference energy 19 MeV. The black line shows the tune shift calculated from the transfer map using normal form analysis; the blue circles show the tune shift from NAFF of tracking results (using the transfer map in Taylor form); and the red points show the tune shift from NAFF of numerical tracking results (using PYZGOUBI). The discrepancy in the vertical plane between PYZGOUBI and the other codes is likely due to the different magnetic field representation used in PYZGOUBI.

$J_x = 200 \mu\text{m}$. The agreement between PYZGOUBI and the transfer map tracking is better than 10^{-5} and gives confidence in both results. However, for actions smaller than $0.1 \mu\text{m}$ the results from PYZGOUBI appear not to converge. This may be due to numerical precision error when tracking particles close to the closed orbit.

In Fig. 10(b), the vertical tune shift with vertical action J_y obtained from normal form analysis (black line) shows good agreement with the result of the NAFF from tracking using a transfer map: the discrepancy in tune shift is less than 10^{-4} for actions smaller than $100 \mu\text{m}$. However, there is a large discrepancy in the variation of the vertical tune with amplitude between tracking in PYZGOUBI (red points) and tracking using the transfer map (blue circles). The discrepancy may be explained by the difference in the description of the magnetic field between the two codes. In PYZGOUBI, the magnetic field throughout the 3D coordinate

space is derived from a 2D field map, in which the field values are given in the median horizontal plane. As explained in Sec. II B, obtaining magnetic field values by extrapolation from values given on a single plane (or, equally, by extrapolation outside a surface from values given on that surface) can be prone to error. On the other hand, extrapolation within a bounded volume from values given on that boundary leads to values in which the error is limited by the error on the boundary: this is the technique used in constructing the transfer map.

To investigate further the discrepancy in vertical tune shift with amplitude between PYZGOUBI and the transfer map, we compare the field observed by the particle as it passes through the vertically focusing quadrupole in each of three different codes: OPERA (which was used to compute the field from the original magnet design); PYZGOUBI (which extrapolates the field from data on the midplane of the magnet); and our transfer map tracking code (which computes the field within a boundary defined by a cylinder inscribed through the magnet). The vertical deflection of a particle is given by the integral of the horizontal field component along the trajectory of the particle. Therefore, we focus on the horizontal component of the field as a function of vertical position, and express the field at the longitudinal midpoint of the magnet as a power series in the vertical coordinate:

$$B_x(y) = \sum_{k=0}^{\infty} c_k y^k, \quad (22)$$

where

$$c_k = \frac{1}{k!} \left. \frac{\partial^k B_x}{\partial y^k} \right|_{y=0}. \quad (23)$$

By fitting a polynomial to the field B_x as a function of y in each of our three codes, we obtain for each code the coefficients c_k of the polynomial. The results are shown in Fig. 11. The coefficient c_0 is related to the closed orbit: the value of c_0 is small in all three codes. This is consistent with expectations: for EMMA, we expect the closed orbit to lie on the midplane. The first-order (quadrupole) coefficient c_1 is related to the tune: there is good agreement in the value for this coefficient for all three codes. This is consistent with the agreement we found previously, for the tunes obtained from numerical tracking in PYZGOUBI, and the tunes obtained from the transfer map. The vertical tune shift with amplitude is related to the third-order (octupole) coefficient, c_3 . We see that OPERA and the transfer map code give good agreement on the value for c_3 ; however, PYZGOUBI gives a different value, and one that is smaller by about 5 orders of magnitude. This is consistent with the much smaller vertical tune shift with amplitude found from tracking in PYZGOUBI, compared to tracking with the transfer map.

We conclude that analysis of the dynamics using the map methods described above can give reliable information on aspects of the nonlinear dynamics, in particular, the

FIG. 11. Value of the coefficients of a polynomial fitted to the horizontal component of the magnetic field as a function of vertical position, for PYZGOUBI (red crosses), transfer map (black circles), and OPERA simulations (blue triangles). Results agree for the first order, corresponding to the quadrupole component of the field. The higher-order coefficients obtained from PYZGOUBI are 5 orders of magnitude smaller than the corresponding coefficients in the case of the transfer map and OPERA. There is a much smaller discrepancy between the transfer map and the OPERA coefficients. We conclude that transfer map represents more accurately than PYZGOUBI the nonlinearities in the magnetic field; this can explain the discrepancy in nonlinear dynamics observed in Fig. 10(b).

tune shift with amplitude, at least in the case when the nonlinear effects are small. However, it is essential to use an accurate description of the field that determines the dynamics. It is possible that the transfer map techniques would also provide reliable information on the nonlinear dynamics where strongly nonlinear fields are present, for example, in the case of a scaling FFAG [39]. However, this is outside the scope of this paper.

IV. USE OF TRANSFER MAPS FOR FITTING A MODEL TO EXPERIMENTAL DATA

A. Interpolation of transfer maps

Having studied the dynamics in a given configuration of the lattice, it is interesting to find a way to extend the use of the transfer map approach, to allow a range of different machine configurations to be easily modeled and compared. During the commissioning of the machine a real-time simulation of the effect of a change in the lattice configuration is extremely useful. When using a hard-edge model for the magnets, it is trivial to study any lattice by directly adjusting the parameters and performing the tracking study. The use of a field map is a significant improvement in the lattice description; but then, in principle, a different field map is needed for each configuration. This would mean that an accurate OPERA model would have to be solved for each new lattice, requiring considerable computer time. Using the superposition of the fields generated by the magnets, taking into account the presence

of the other magnet as explained previously, the construction of the field map for any given machine configuration (specified by particular values of the magnet strengths and positions) can be performed easily. However, the dynamical properties of the lattice need to be determined either by particle tracking, or by construction of a transfer map. In either case, the necessary computations are not particularly lengthy, but still take some minutes to perform.

In the case of numerical tracking, little can be done to make the process more efficient. The only way to determine properties such as the tunes and time of flight is to carry out numerical integration of the equations of motion for each new lattice configuration. However, when using transfer maps, we can consider constructing a grid of “reference” maps, corresponding to a set of lattice configurations, and then obtaining the transfer map for any desired new configuration by interpolation between the reference maps. The interpolation can be performed essentially instantly, and the dynamical properties of the new configuration can then be obtained immediately from the map, without the need for any tracking.

In developing a technique based on interpolating transfer maps, questions arise as to the density of reference maps needed, and the best way to perform the interpolation, to achieve maps that are sufficiently reliable in predicting the behavior of the machine. The interpolation technique can be tested by comparing the dynamical properties predicted from an interpolated map, with the properties predicted by a map computed for the same lattice configuration. If the differences between the predicted properties are small compared to experimental error, then the interpolation technique could be useful in practice.

The number of reference maps required is (so far) found by trial and error. For the interpolation method itself, there are at least two possible approaches. The most simple and direct is to interpolate the individual coefficients of the Taylor map. However, without any explicit constraint, the precise relationship between the coefficients that ensures the symplecticity of the map is lost in the interpolation: the result is a map with potentially large symplectic error. We have found that this technique produces maps that not only have a large symplectic error, but that are also very unreliable in their prediction for the dynamical properties of different lattices. One possible reason for this is that the overall accuracy of the interpolated map reduces with the number of quantities being interpolated; in the case of a high-order Taylor map for six dynamical variables, the number of interpolated quantities (power series coefficients) can be extremely large.

We can hope to improve the reliability of the interpolation by reducing the number of quantities that need to be interpolated. One way to do this is to use a representation of the map in which only the minimum number of coefficients are used to specify the map. For a symplectic map, there are several possibilities, including Lie generators,

and mixed variable generating functions (MVGFs). The symplectic property is exploited in the sense that symplecticity imposes a strong constraint on the map. In the Taylor series representation, the coefficients are related in specific ways, although the relationships are usually too complex to write down explicitly. However, by converting the map into an MVGF (for example), we construct a representation in which the coefficients involved are all (in principle) independent, and their number is greatly reduced compared to a Taylor series representation. The price to pay is that some work is needed to convert the various maps between various representations. However, standard techniques exist to convert a (symplectic to some order) Taylor map into an MVGF, and the conversion only needs to be done once for each reference map. Using COSY, it is possible to construct an MVGF that reproduces any transfer map symplectic to order n from the map expressed in the form of a Taylor series [27,36]. An MVGF is not in explicit form (i.e. it cannot be used directly for particle tracking); however, a map represented in the form of an MVGF is necessarily symplectic. The interpolation is performed on the MVGFs; the interpolated map can then readily be converted back into Taylor form using a DA code (such as COSY), by “solving” the map represented by the MVGF, using DA variables to represent the dynamical variables.

We note that MVGFs have long been used as representations for transfer maps, or to “symplectify” Taylor maps (see, for example, [24,29,54–58]). However, the motivation (and application) has generally been to perform symplectic tracking, so as to be able to make accurate predictions of dynamic aperture in proton storage rings. Here, we use MVGFs in a slightly different way: by interpolating between maps in the form of MVGFs rather than in the form of Taylor series, we improve the accuracy of the interpolated map. Since the beam in an FFAG is accelerated over a relatively small number of turns, dynamic aperture is not generally an issue for an FFAG, so the need for the map to be symplectic arises from other concerns. In the present case, using an implicitly symplectic representation of the map allows us to improve the accuracy of the map obtained by interpolating between maps representing different machine configurations.

It would be possible, in principle, to include the lattice parameters (such as quadrupole strengths and positions) as DA variables in COSY, in addition to the dynamical variables. Then, one could produce a single “map” in Taylor form, which would include dependence on the lattice configuration through appropriate parameters, which would be specified before evaluating the map. This approach would require the magnetic field map to be represented in an appropriate form: the most convenient representation would be one in which the field itself includes dependence on parameters specifying the lattice configuration. However, as explained in Sec. II A, although the field maps for different quadrupole strengths can be obtained by

scaling and superposition of the maps for each magnet, since the yoke of each magnet in EMMA affects the field of the other magnet, a different field map needs to be computed for each configuration in which the magnets have different relative horizontal positions. Including the positions of the magnets as parameters in a map generated in COSY is therefore not trivial, and we have not carried this through, though it may be an interesting topic for future study.

In our investigations, we used a grid of 300 reference lattices (i.e. 300 different configurations). Finding the closed orbit and time of flight by numerical integration in PYZGOUBI for 300 reference lattices, at nine different energies, takes a little over 4 hours. Constructing the transfer maps for the same configurations takes about 5 hours; dynamical properties such as the closed orbit and time of flight can then be obtained for the full set of configurations in less than a minute. In this respect, it appears that the computational time required for the transfer map approach is longer than that for numerical integration, for the same task. However, to compute the time of flight at an additional reference energy (for example) will take only a few seconds using the transfer maps, whereas an additional half an hour would be needed using PYZGOUBI. Furthermore, computing additional dynamical quantities, such as the betatron tunes, for all 300 reference lattices and 9 reference energies would require several hours more computation time in PYZGOUBI, whereas the same information can be directly extracted from the transfer maps already computed, within a few seconds.

For each reference lattice, an MVGF was computed from the Taylor map obtained by integration. We found that for all lattice configurations, an MVGF of the second kind (using Goldstein’s nomenclature [59]) could be constructed. An MVGF of the second kind is a function of the old coordinates and the new momenta:

$$F_2 = F_2(q_i, p_f, t), \quad (24)$$

where q_i are the old coordinates and q_f are the new coordinates, p_i and p_f are their respective conjugate momenta. The old and new variables are related by

$$p_i = \frac{\partial F_2}{\partial q_i}, \quad q_f = \frac{\partial F_2}{\partial p_f}. \quad (25)$$

A simple linear interpolation algorithm was used to find the MVGF for a specified lattice “in between” the reference lattices.

B. Symplectic error

The effect of interpolating the MVGF rather than the coefficients of the Taylor map is illustrated clearly by comparing the symplectic error in the maps. For purposes of making a comparison, we characterize the symplectic error as follows. First, we note that a symplectic map is defined by the condition

FIG. 12. Terms in the matrix E characterizing the symplectic error, with increasing order of the terms, for particular values of the dynamical variables. (a) Map obtained by interpolation of the coefficients of Taylor maps, with dynamical variables given the value 0.001. (b) Map obtained by interpolation of MVGFs, with dynamical variables given the value 0.01.

$$J^T S J = S, \quad (26)$$

where J is the Jacobian of the map, and S is a block-diagonal matrix, with block diagonals:

$$S_2 = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (27)$$

Then, we construct the matrix E , given by

$$E = J^T S J - S. \quad (28)$$

For a transfer map expressed in Taylor form, each term of the Jacobian is a polynomial in the dynamical variables. Thus, the components of E are also polynomials. To evaluate numerically the components of E , we must replace the variables by some values; these values can be chosen arbitrarily. In general, we find that for sufficiently large values of the dynamical variables, the numerical values of the terms in each component of E increase with increasing order of the term. We can characterize the symplectic error in a map, as the maximum value of the dynamical variables

which, when substituted (for all variables) into each term in each component in E , does not lead to an increasing value with increasing order of the term.

Figure 12 shows the results of the analysis of the symplectic error in this way for two cases: a map obtained by direct interpolation of the coefficients of the Taylor maps; and a map obtained for the same configuration, by interpolation of the MVGF. In the former case, the largest value for the dynamical variables that does not lead to an increase in the terms in E with increasing order is 0.001. This indicates that we may expect to see effects of the symplectic error for values of x or y larger than 1 mm (for example). On the other hand, when constructing a map using interpolation of the MVGF, then a value of 0.01 can be substituted for the dynamical variables, before the terms in E start to increase with increasing order. Furthermore, the values of the terms in E at this point are 3 orders of magnitude smaller than for the map constructed by interpolation of the Taylor maps with value 0.001 for the dynamical variables. This indicates that the symplectic error in the map constructed by interpolation of the Taylor map coefficients is significantly larger than for the map constructed from interpolation of the MVGFs. In fact, since a map represented by an MVGF is necessarily symplectic, the symplectic error must come from numerical errors in computing the coefficients, and from truncation of the map, when the map is converted to Taylor form.

C. Accuracy of interpolated maps

To test the interpolation technique, we compare the beam dynamics in an interpolated map, with the beam dynamics in a map computed directly for the same specified configuration. A four-dimensional grid was first constructed for a range of different lattice configurations in EMMA (the four dimensions corresponding to the strengths and positions of the two magnets in one cell). Transfer maps with 14 MeV reference kinetic energy were computed for 300 lattice configurations: F and D magnet strengths, A_f and A_d , were varied from 94% to 110% of the nominal magnet strength in 4% steps; the transverse positions of the axes of the D magnet and the F magnet with respect to the reference polygon (x_d and x_f respectively, in Fig. 1) were varied from 32 to 38 mm and from 8 to 12 mm, respectively, in 2 mm steps.

A particular configuration, midway between four grid points was then selected, and the map for this configuration constructed by interpolation of the MVGFs, as described above. The map for the selected configuration was also computed directly, by symplectic integration through the fields using COSY.

Figure 13 shows the horizontal and vertical tunes for various energies, obtained from a directly computed Taylor map and from a Taylor map derived from an interpolated MVGF. There is good agreement between the tunes obtained from the directly computed map, and from the

FIG. 13. Horizontal (blue) and vertical (red) tune per cell at various energies, found from a directly computed map, and from an interpolated MVGF map. The reference energy is 14 MeV. There is some small discrepancy in the results from different maps, but it is below the measurement precision of 0.01.

interpolated MVGF: the discrepancy is below the measurement precision. Figure 14 shows the time of flight as a function of energy, obtained from a directly computed Taylor map, and from a Taylor map derived from an interpolated MVGF. Again, there is good agreement (better than measurement precision) between the two maps, indicating the accuracy of the interpolated map. Also shown, for comparison, is the time-of-flight curve for a map obtained by interpolating directly the Taylor map coefficients (i.e. without use of MVGFs). In this case, the discrepancy with the

FIG. 14. Variation of time of flight with energy, found from a directly computed map, a map obtained by interpolating the Taylor map coefficients, and a map obtained from interpolation between the MVGF coefficients. The reference energy for all the maps is 14 MeV. We observe good agreement between the directly computed map and the map obtained by interpolating the MVGF.

FIG. 15. Betatron motion with multiple transfer maps of order eight, for five energy deviations with 14 MeV reference energy. For the computed map (black) and the map derived from an interpolated MVGF (red), effects of a symplectic error are apparent for energy deviations of 3 MeV (20%).

directly computed map is significantly larger. As well as enforcing symplecticity, use of the MVGF interpolation also appears to improve the accuracy of the interpolated map.

Figure 15 shows the horizontal phase space obtained by tracking using the directly computed Taylor map, and a map from the interpolated MVGF. There is good agreement between the results for a range of energies (a reference energy of 14 MeV was used in computing the maps, and the energy variation achieved by setting the energy deviation up to 20%). In both cases, nonsymplectic behavior (a steady “spiralling out” of the particle trajectory in phase space) is visible for energy deviations larger than about 15% (2 MeV deviation from 14 MeV reference energy).

V. COMPARISON WITH EXPERIMENTAL DATA

The highly compact lattice in EMMA makes injection (and extraction) particularly challenging. For commissioning, therefore, a baseline lattice configuration was selected, in which it was expected to be relatively straightforward to achieve injection. Early experiments [5] were performed using a fixed injection energy of 12 MeV. The dynamical properties of the lattice at other energies were studied by varying the magnet strengths: reducing the magnet strengths by 10%, for example, gave the same behavior as would be achieved by increasing the beam momentum by 10%.

In practice, it was found that injection was most easily achieved using a lattice slightly different from the nominal baseline. We specify the lattice configuration by giving the magnet strengths relative to the baseline lattice, and the absolute magnet positions. Thus, the lattice used for the measurements reported here had the horizontal focusing

quad in each cell at 105.6% of nominal strength, the horizontally defocusing quad at 100% of nominal strength, and these magnets at positions $x_f = 9.51$ mm and $x_d = 36.05$ mm (see Fig. 1). This lattice, which we refer to as the E_1 lattice, is then specified by the notation (105.6, 100, 9.51, 36.05).

A transfer map for the E_1 lattice was constructed using MVGF interpolation. Figure 16(a) shows the comparison between measurements and simulations of the time of flight as a function of energy for this lattice configuration. We observe that the measured lattice has a minimum time

FIG. 16. (a) Comparison between measurements and simulations of the time-of-flight variation with energy for the baseline lattice configuration. We observe that the measured lattice has a minimum time of flight of 55.275 ns at 15 MeV whereas the model lattice (with the same nominal magnet strengths and positions as used in the machine) has a minimum time of flight of 55.225 ns at 17 MeV. (b) From the time-of-flight measurements (blue points with error bars) four constraints (red triangles) are defined for the optimization routine used to find a model lattice configuration with the same dynamical behavior as observed in the machine. The time of flight for a range of energies between 10 and 20 MeV is then computed for the fitted lattice (black line with circles) found by the optimization routine. The time of flight in the fitted lattice is in good agreement with the time-of-flight measurements.

of flight of 55.275 ns at 15 MeV (“equivalent energy”) whereas the simulated lattice has a minimum time of flight of 55.225 ns at 17 MeV. There is clearly some significant difference between the machine and the model.

To investigate the difference between the model and the machine, the grid of lattice configurations was used to find a model that corresponded to the measurements. To do this, an optimization routine (based on the Nelder-Mead method, implemented in MATHEMATICA [46,60,61]) was used to look for the model lattice configuration that most closely matched the measured time of flight at different energies. For each lattice configuration of the grid, we used the transfer map to compute the time of flight on the closed orbit from 10 to 20 MeV, varying the energy deviation. In many lattice configurations, especially close to the edges of the grid, no closed orbit could be found for energies higher than 18 MeV; therefore the time of flight could only be interpolated between 10 and 18 MeV. Once the configuration parameters for the lattice best fitting the time-of-flight curve had been found, the transfer map was constructed by interpolation of the MVGF.

Figure 16(b) shows a comparison between measurements and simulations from the lattice configuration found by the optimization routine (the constraints used for the fit are shown as red triangles). The fitted lattice found by the optimization routine, which we refer to as S_1 , has parameters (in the same notation used above for E_1): (94.30, 100.02, 11.56, 34.00). The times of flight for closed orbits in the energy range between 10 and 20 MeV were computed for this lattice (black line with circles). It shows good agreement with the measurements (blue points with error bars) with a maximum discrepancy of about 20 ps at 10.5 MeV. All the other measurements agree within less than 10 ps.

Lattices S_1 and E_1 are then compared in terms of tune per cell. The experimental values of the tune are obtained from the orbit transverse position [given by beam position monitors (BPMs) around the ring] by the NAFF method. We observe in Fig. 17(a) that there is relatively good agreement for the horizontal tune (in blue) for energies below 16 MeV. The simulated vertical tune (red circles) is significantly larger than the measurement, especially for large energy. Note that the simulated tune is computed from a single transfer map with a reference energy of 14 MeV.

Comparing Fig. 17(a) with Fig. 7, we observe that computing the tune from maps with different reference energies gives smaller values for the tune at high energy, compared with the tune obtained by varying the energy deviation in a single map. This may explain part of the discrepancy. In order to interpolate an MVGF with a different reference energy, a new grid of transfer maps must be computed for each energy. Instead, we can derive the analytical representation of the magnetic field for the given configuration and directly compute the transfer maps with different reference energies. The level of agreement between measurements

and multiple computed map simulations is better than 0.01 for the horizontal tune [see Fig. 17(b)]. The simulated vertical tune is still larger than the measurement; however, the discrepancy at high energy is significantly reduced.

We conclude that the lattice found by the optimization routine to match the measurement agrees also with the measurements of the tunes at various energies. This is an interesting result since it gives confidence in the use of the optimization routine to identify a lattice that matches the behavior in the machine in more than just the characteristics used as constraints in the optimization.

The simulation lattice S_1 that matches the time-of-flight measurements has significantly different configuration parameters compared to the experimental lattice E_1 . Although simulations and measurements do not agree in

(a)

(b)

FIG. 17. Tunes per cell in the lattice S_1 (circles) are compared with the measured tunes from the lattice E_1 (points). In (a), the tunes in S_1 are obtained from a single transfer map (interpolated MVGF) with reference energy 14 MeV. There is a relatively good agreement for the horizontal tune (in blue) for energies smaller than 16 MeV. The simulated vertical tune (red) is significantly larger than the measurement, especially for large energy. In (b), the tunes are computed from a set of transfer maps for the same lattice configuration, but with different reference energies, and zero energy deviation. The agreement with the measurements at higher energy is improved.

absolute terms, it is interesting to compare the response to changes in the lattice predicted by the model with the response to changes measured in the machine. To do so, we measured the time of flight for various energies for three transverse positions of both magnets: both moved 1 mm inward relative to E_1 (lattice E_2 , -1 mm offset), both moved 2 mm inward (lattice E_3 , -2 mm offset), and both moved 3 mm inward (lattice E_4 , -3 mm offset). The magnet strengths are kept constant; thus, for instance, the lattice E_2 is (105.6, 100, 8.51, 35.05). The time-of-flight measurements for these lattices are represented by points with error bars in Fig. 18. We then apply the corresponding magnet movements to the lattice S_1 and obtain the lattices S_2 , S_3 , and S_4 ; for instance, lattice S_2 is (94.30, 100.02, 10.56, 33.00). For each corresponding change in positions of the magnets, the experimental and simulated results are plotted in the same color. Note that lattice S_4 is actually outside the grid of reference lattices ($x_d < 32$ mm), and hence the time of flight as a function of energy needs to be found by extrapolation, rather than interpolation. We observe that the agreement between measurement and simulation is within 20 ps, apart from some specific measurements (for which the actual kinetic energy of the injected might not have been exactly 12 MeV, because of rf phase jitter and electron gun instabilities in ALICE). There is good agreement between model and measurement on the position (in energy and time of flight) of the minimum of the time-of-flight curve for each lattice configuration: the position of the minimum is important for determining the optimum rf parameters for achieving acceleration in EMMA.

We can then conclude that simulations can be used to predict the effect on the time of flight, at least when moving both magnets by the same amount. Predictions from the

FIG. 18. Experimental measurements of the time of flight (TOF) per turn are compared with simulation results for four different lattice configurations. A model is initially fitted to TOF measurements in a single lattice configuration at four different energies; then the models for new configurations are constructed by applying the same changes to the fitted model, as are applied in the machine.

fitted model could be used to optimize acceleration. Further investigations should explore the ability of the fitted model to predict the behavior of the machine under a range of larger and more complicated configuration changes.

Finally, we note that the tune per cell measurements are obtained from the transverse position of the beam in each cell recorded over many successive turns. At the stage of commissioning during which the measurements reported here were collected, the transverse position of the orbit was measured by directly connecting the BPM button signals to an oscilloscope in the control room; measuring the orbit around the entire ring required therefore a significant time (for recabling, and collecting and processing the data) and it was not done for the other experimental lattices described here. More recently, all the BPMs have been connected to an EPICS interface [62], making the measurement of orbit position and tunes almost instantaneous. This should enable a more rigorous and complete evaluation of the optimization routine to be made in the future.

VI. SUMMARY AND CONCLUSIONS

The techniques presented in this study consist of the application of four independent tasks: first, construction of an accurate model of a magnetic field in a given section of accelerator beam line; second, derivation of an accurate analytical representation of the given magnetic field; third, use of an integrator in a differential algebra code to construct a transfer map in Taylor form; and finally, use of a range of analytical techniques (including, for example, computation of the Lie generator for a given Taylor map, and conversion of Taylor maps to a mixed variable generating function representation, and *vice versa*) to extract significant features of the dynamics. The use of transfer maps has enabled the development of an efficient tool for exploring the effects of changing the lattice configuration, including the fitting of a model of a lattice to experimental measurements.

The nonscaling FFAG accelerator EMMA was chosen as a case study for three main reasons. First, the magnets have fringe fields that have a strong influence on the dynamics. Accurate simulations of the dynamics depend on accurate models for the magnets. Second, the flexibility of the lattice, with 4 degrees of freedom, requires also much flexibility in the codes used to model the dynamics. Exploring the effect of changes in lattice configuration, and matching a configuration in the model to measured dynamical characteristics, requires efficient tools capable of including detailed field models. Third, the large excursion of particle trajectories with respect to a given straight reference trajectory tests the capability of mapping techniques that necessarily use the paraxial approximation.

The integrator we used to construct the transfer map was chosen partly on the grounds that it is symplectic. At first sight, the symplecticity of the integrator does not appear to

be a major consideration: in a machine such as EMMA, particles are accelerated over the course of a relatively small number of turns, during which time the effects of some symplectic error in an otherwise reasonably accurate map should be small; and in any case, a symplectic error is introduced by the truncation of a Taylor map during the integration using a differential algebra code. However, use of a symplectic integrator was found to be convenient for two reasons. First, the Taylor map produced by the integrator is symplectic up to a given order: verification of the symplecticity to this order provides a useful check that the integrator has been implemented correctly. Perhaps more interestingly, we found that enforcing symplecticity in interpolating the transfer map (by use of mixed variable generating functions) provided more accurate results than were obtained by interpolating directly the coefficients in a Taylor map.

We have seen that transfer maps provide a convenient tool for modeling the dynamics in a nonscaling FFAG such as EMMA, with some advantages over purely numerical modeling techniques. In the case of EMMA, the dynamics are essentially linear, and it is possible to describe the dynamics over the full energy range (for a given lattice configuration) using three maps, each with a different reference energy. Nonlinear effects, such as tune shift with amplitude, are small. However, the techniques that we have developed and described here should be equally applicable to the case, for example, of a scaling FFAG, where the shapes of the magnetic fields lead to much stronger nonlinear effects. Furthermore, techniques using transfer maps may provide powerful tools for the design of such accelerators, given the capability of these techniques to allow efficient analysis of a range of important dynamical effects.

ACKNOWLEDGMENTS

We wish to thank our colleagues at the EMMA commissioning team for helpful discussions and technical support. Also, we are grateful to R. Nilavalan for his support of this work. Finally, we would like to express our particular gratitude to S. Machida, D. Kelliher, and B. Muratori for their helpful comments on this paper. This work was supported by the Science and Technology Facilities Council, UK.

APPENDIX: CONVERSION OF MAGNETIC FIELD MODE COEFFICIENTS FROM CYLINDRICAL TO CARTESIAN BASIS

We consider the case that the radial component of the magnetic field is given by

$$B_\rho = \sum_{m,n} a_{mn} I'_m(nk_z \rho) \sin(m\phi) \cos(nk_z z), \quad (\text{A1})$$

and the vertical field component in the same field is given by

$$B_y = \sum_{m,n} c_{mn} \cos(mk_x x) \cosh(k_y y) \cos(nk_z z), \quad (\text{A2})$$

where

$$k_y^2 = m^2 k_x^2 + n^2 k_z^2. \quad (\text{A3})$$

Given a set of cylindrical mode coefficients a_{mn} , our objective is to find the corresponding Cartesian mode coefficients c_{mn} .

Since the fields are the same everywhere, and are determined by the coefficients, we can choose a particular point in space. If we can find the coefficients c_{mn} by considering the fields just at that particular point, then we have solved the problem.

Let us choose the point $z = x = 0$. Since $x = 0$, it follows that $\phi = \pi/2$. Let us assume that the symmetry of the field is such that only odd values of m appear. The expressions (A1) and (A2) become at this point

$$B_\rho = \sum_{m,n} a_{mn} I'_m(nk_z \rho) (-1)^{(m-1)/2}, \quad (\text{A4})$$

$$B_y = \sum_{m,n} c_{mn} \cosh(k_y y). \quad (\text{A5})$$

Since the radial field is along the y axis at this point, then we can say that $y = \rho$, and $B_y = B_\rho$. Therefore,

$$\sum_{m,n} a_{mn} I'_m(nk_z \rho) (-1)^{(m-1)/2} = \sum_{m,n} c_{mn} \cosh(k_y \rho). \quad (\text{A6})$$

Now we assume a linear relationship between the coefficients:

$$c_{mn} = f_{mm'}^{(n)} a_{m'n}. \quad (\text{A7})$$

Our goal is to find the components $f_{mm'}^{(n)}$, which will allow us to transform from the cylindrical to the Cartesian basis.

Substituting Eq. (A7) into Eq. (A6) gives

$$\sum_{m,n} a_{mn} I'_m(nk_z \rho) (-1)^{(m-1)/2} = \sum_{m',m,n} f_{mm'}^{(n)} a_{m'n} \cosh(k_y \rho). \quad (\text{A8})$$

This equation is satisfied (a sufficient but not necessary condition) if each term in the summation over n on one side is the same as the corresponding term on the right. Then, we can simply drop the summation over n :

$$\sum_m a_{mn} I'_m(nk_z \rho) (-1)^{(m-1)/2} = \sum_{m',m} f_{mm'}^{(n)} a_{m'n} \cosh(k_y \rho). \quad (\text{A9})$$

Now we make use of the series expansion for the derivative of the modified Bessel function:

$$I'_m(\xi) = \sum_{l=0}^{\infty} \frac{(2l+m)\xi^{2l+m-1}}{2^{2l+m} l! (l+m)!} \quad (\text{A10})$$

and the hyperbolic cosine:

$$\cosh(\xi) = \sum_{l=0}^{\infty} \frac{\xi^{2l}}{(2l)!}. \quad (\text{A11})$$

Substituting Eqs. (A10) and (A11) into Eq. (A9) gives

$$\begin{aligned} \sum_l \sum_m a_{mn} (-1)^{(m-1)/2} \frac{(2l+m)(nk_z \rho)^{2l+m-1}}{2^{2l+m} l! (l+m)!} \\ = \sum_l \sum_{m',m} f_{mm'}^{(n)} a_{m'n} \frac{(k_y \rho)^{2l}}{(2l)!}. \end{aligned} \quad (\text{A12})$$

Equation (A12) must be true for all values of ρ . Therefore, the coefficients of equal powers of ρ on either side must be equal. Let us consider the coefficients of $\rho^{2l'}$. On the right-hand side of Eq. (A12), this is simply the term in the summation over l with $l = l'$. On the left, this is the term in the summation over l with

$$2l + m - 1 = 2l', \quad \therefore \quad l = l' - \frac{m-1}{2}. \quad (\text{A13})$$

Hence, comparing coefficients of powers of ρ gives

$$\begin{aligned} \sum_m a_{mn} (-1)^{(m-1)/2} \frac{(2l'+1)(nk_z)^{2l'}}{2^{2l'+1} (l' - \frac{m-1}{2})! (l' + \frac{m+1}{2})!} \\ = \sum_{m',m} f_{mm'}^{(n)} a_{m'n} \frac{k_y^{2l'}}{(2l')!}. \end{aligned} \quad (\text{A14})$$

Let us relabel l' as l on both sides, and interchange the labels m and m' on the right:

$$\begin{aligned} \sum_m a_{mn} (-1)^{(m-1)/2} \frac{(2l+1)(nk_z)^{2l}}{2^{2l+1} (l - \frac{m-1}{2})! (l + \frac{m+1}{2})!} \\ = \sum_{m',m} f_{m'm}^{(n)} a_{mn} \frac{k_y^{2l}}{(2l)!}. \end{aligned} \quad (\text{A15})$$

As before we can equate corresponding terms in the summation; in this case, we take the summation over m :

$$(-1)^{(m-1)/2} \frac{(2l+1)(nk_z)^{2l}}{2^{2l+1} (l - \frac{m-1}{2})! (l + \frac{m+1}{2})!} = \sum_{m'} f_{m'm}^{(n)} \frac{k_y^{2l}}{(2l)!}. \quad (\text{A16})$$

Since there is no summation over l , this becomes

$$(-1)^{(m-1)/2} \frac{(2l+1)!(nk_z)^{2l}}{2^{2l+1} (l - \frac{m-1}{2})! (l + \frac{m+1}{2})!} = \sum_{m'} k_y^{2l} f_{m'm}^{(n)}. \quad (\text{A17})$$

Now we observe we can write this in matrix form as

$$v_{lm}^{(n)} = \sum_{m'} \bar{k}_{lm'}^{(n)} f_{m'm}^{(n)}, \quad (\text{A18})$$

where we define the matrix $v^{(n)}$ with components

$$v_{lm}^{(n)} = (-1)^{(m-1)/2} \frac{(2l+1)!(nk_z)^{2l}}{2^{2l+1} (l - \frac{m-1}{2})! (l + \frac{m+1}{2})!}, \quad (\text{A19})$$

and the matrix $\bar{k}^{(n)}$ with components

$$\bar{k}_{lm'}^{(n)} = k_y^{2l} = (m'^2 k_x^2 + n^2 k_z^2)^l. \quad (\text{A20})$$

Formally, the matrices have an infinite number of components; however, we can reduce them to finite-sized square matrices by making appropriate truncations in l , m , and m' . We can then solve Eq. (A18) for the required components $f_{mm'}^{(n)}$:

$$f_{m'm}^{(n)} = \sum_l (\bar{k}^{(n)})_{m'l}^{-1} v_{lm}^{(n)}, \quad (\text{A21})$$

or

$$f_{mm'}^{(n)} = \sum_l (\bar{k}^{(n)})_{ml}^{-1} v_{lm'}^{(n)}. \quad (\text{A22})$$

Since Eqs. (A19) and (A20) give us explicit constructions for $v^{(n)}$ and $\bar{k}^{(n)}$ (in terms of the known quantity k_z , and the “freely chosen” quantity k_x), we can solve Eq. (A22) to find the components $f_{mm'}^{(n)}$; and then we can use Eq. (A7) to find the mode coefficients in the Cartesian basis from the mode coefficients in the cylindrical basis.

-
- [1] K. L. Brown, SLAC Report No. 75 (Revision 4, 1982).
 - [2] J. S. Berg, ICFA Beam Dynamics Newsletter **43**, 60 (2007).
 - [3] J. S. Berg, Nucl. Instrum. Methods Phys. Res., Sect. A **596**, 276 (2008).
 - [4] R. Edgecock *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **624**, 1 (2010).
 - [5] S. Machida *et al.*, Nature Phys. **8**, 243 (2012).
 - [6] S. Machida and D. Kelliher, Phys. Rev. ST Accel. Beams **10**, 114001 (2007).
 - [7] E. Keil, A. M. Sessler, and D. Trbojevic, Phys. Rev. ST Accel. Beams **10**, 054701 (2007).
 - [8] D. Trbojevic, E. Keil, and A. M. Sessler, ICFA Beam Dynamics Newsletter **43**, 115 (2007).
 - [9] C. Johnstone and S. Koscielniak, ICFA Beam Dynamics Newsletter **43**, 125 (2007).
 - [10] S. Machida, Phys. Rev. ST Accel. Beams **11**, 094003 (2008).
 - [11] S. Machida, Phys. Rev. ST Accel. Beams **9**, 104002 (2006).
 - [12] J. S. Berg, Phys. Rev. ST Accel. Beams **9**, 034001 (2006).
 - [13] D. Trbojevic, E. D. Courant, and M. Blaskiewicz, Phys. Rev. ST Accel. Beams **8**, 050101 (2005).
 - [14] S. L. Sheehy, K. J. Peach, H. Witte, D. Kelliher, and S. Machida, Phys. Rev. ST Accel. Beams **13**, 040101 (2010).
 - [15] N. Marks, B. J. A. Shepherd, B. Leigh, F. Goldie, and M. J. Crawley, in *Proceedings of the 11th European Particle Accelerator Conference, Genoa, 2008* (EPS-AG, Genoa, Italy, 2008).
 - [16] Vector Field OPERA 3D, <http://www.vectorfield.com/>.
 - [17] F. Méot and S. Valero, CEA DSM Report No. DAPNIA/SEA-97-13, 1997.
 - [18] F. Méot, ICFA Beam Dynamics Newsletter **43**, 44 (2007).
 - [19] S. Tygier and D. Kelliher, <http://sourceforge.net/projects/pyzgoubi/>.
 - [20] Y. Giboudot and F. Méot, in *Proceedings of the 23rd Particle Accelerator Conference, Vancouver, Canada, 2009* (IEEE, Piscataway, NJ, 2009).
 - [21] A. W. Chao, T. Sen, Y. T. Yan, and E. Forest, *Proceedings of the 1991 Particle Accelerator Conference, San Francisco, CA, USA* (APS Beam Physics, San Francisco, California, 1991), pp. 336–338.
 - [22] T. Sen, Y. Yan, A. Chao, and E. Forest, Report No. SSCL-497, 1991.
 - [23] R. Kleiss, F. Schmidt, Y. Yan, and F. Zimmerman, DESY HERA 92-01, CERN SL/92-02 (AP), and SSCL-564, 1992.
 - [24] Y. T. Yan, P. J. Channell, M. Li, and M. J. Syphers, in *Proceedings of the 1993 Particle Accelerator Conference, Washington, DC* (IEEE, New York, 1993), pp. 38–40.
 - [25] W. Scandale, in *Proceedings of the 1993 Particle Accelerator Conference, Washington, DC* (Ref. [24]), pp. 16–20.
 - [26] E. Forest, *Beam Dynamics: A New Attitude and Framework* (Harwood Academic Publishers, Amsterdam, 1998).
 - [27] M. Berz, *Modern Map Methods in Particle Beam Physics* (Academic Press, San Diego, 1999).
 - [28] D. T. Abell, E. McIntosh, and F. Schmidt, Phys. Rev. ST Accel. Beams **6**, 064001 (2003).
 - [29] A. J. Dragt, “Lie Methods for Nonlinear Dynamics with Applications to Accelerator Physics,” (2011), <http://www.physics.umd.edu/dsat/dsatliemethods.html>.
 - [30] M. Venturini and A. J. Dragt, Nucl. Instrum. Methods Phys. Res., Sect. A **427**, 387 (1999).
 - [31] A. J. Dragt, P. Roberts, T. J. Stasevich, A. Bodoh-Creed, and P. Walstrom, in *Proceedings of the 2000 International Computational Accelerator Physics Conference, Darmstadt, Germany* (2000).
 - [32] C. E. Mitchell and A. J. Dragt, ICFA Beam Dynamics Newsletter **42**, 65 (2007).
 - [33] M. Berz and E. Forest, Report No. SSC-N-571, 1988.
 - [34] M. Berz, Part. Accel. **24**, 109 (1989).
 - [35] M. Berz, in *Handbook of Accelerator Physics and Engineering*, edited by A. W. Chao and M. Tigner (World Scientific, Singapore, 1999), pp. 82–85.
 - [36] M. Berz *et al.*, COSY Infinity, http://bt.pa.msu.edu/index_cosy.htm.
 - [37] Y. Giboudot, Ph.D. thesis, Brunel University, 2011.
 - [38] Y. K. Wu, E. Forest, and D. S. Robin, Phys. Rev. E **68**, 046502 (2003).
 - [39] Y. Mori and M. Aiba, ICFA Beam Dynamics Newsletter **43**, 27 (2007).
 - [40] A. J. Dragt and E. Forest, J. Math. Phys. (N.Y.) **24**, 2734 (1983).
 - [41] A. J. Dragt and J. M. Finn, J. Math. Phys. (N.Y.) **20**, 2649 (1979).
 - [42] E. Forest, Report No. SSC-29, 1985.
 - [43] A. Bazzani, P. Mazzanti, G. Servizi, and G. Turchetti, Nuovo Cimento B **102**, 51 (1988).
 - [44] E. Forest, M. Berz, and J. Irwin, Part. Accel. **24**, 91 (1989).
 - [45] M. Berz, AIP Conf. Proc. **249**, 456 (1991).
 - [46] <http://www.wolfram.com/>.
 - [47] J. Laskar, Icarus **88**, 266 (1990).
 - [48] J. Laskar, Physica (Amsterdam) **67D**, 257 (1993).

- [49] S. Dumas and J. Laskar, *Phys. Rev. Lett.* **70**, 2975 (1993).
- [50] R. Bartolini, M. Giovannozzi, W. Scandale, A. Bazzani, and E. Todesco, in *Proceedings of EPAC'96*, Sitges, Spain (1996).
- [51] J. Laskar, in *Proceedings of NATO Advance Institute 533 3DHAM95*, S'Agaro, Spain, 1995, edited by C. Simó (1999).
- [52] Y. Papaphilippou, in *Proceedings of the 18th Particle Accelerator Conference, New York, 1999* (IEEE, New York, 1999), pp. 1554–1556.
- [53] D. Robin, C. Steier, J. Laskar, and L. Nadolski, *Phys. Rev. Lett.* **85**, 558 (2000).
- [54] Y. Yan, P. Channell, and M. Syphers, Report No. SSCL-Preprint-157, 1992.
- [55] J. S. Berg, R. L. Warnock, R. D. Ruth, and E. Forest, *Phys. Rev. E* **49**, 722 (1994).
- [56] R. L. Warnock and J. S. Berg, Report No. SLAC-PUB-7464, 1997.
- [57] B. Erdélyi and M. Berz, *Phys. Rev. Lett.* **87**, 114302 (2001).
- [58] E. Forest, *J. Phys. A* **39**, 5321 (2006).
- [59] H. Goldstein, C. P. Poole, and J. L. Safko, *Classical Mechanics* (Pearson Education, 2001), 3rd ed.
- [60] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms* (MIT Press and McGraw-Hill, Massachusetts, 2001), 2nd ed.
- [61] MATHEMATICA script written by J. Jones and adapted to the processing of experimental results from EMMA.
- [62] J. S. Berg, Y. Giboudot, D. J. Kelliher, S. Machida, F. Méot, B. Shepherd, and S. Tygier, *Proceedings of IPAC'10, Kyoto, Japan* (ICR, Kyoto, 2010), pp. 4325–4327.