APPLICATION OF DYNAMICAL MAPS TO THE FFAG EMMA COMMISSIONING *

Y. Giboudot†, R.Nilavalan, Brunel University, Uxbridge, UK.
A. Wolski, University of Liverpool, Liverpool, UK.

Abstract

The lattice of the Non Scaling FFAG EMMA has four degrees of freedom (strengths and transverse positions of each of the two quadrupoles in each periodic cell). Dynamical maps computed from an analytical representation of the magnetic field may be used to predict the beam dynamics in any configuration of the lattice. An interpolation technique using a mixed variable generating function representation for the map provides an efficient way to generate the map for any required lattice configuration, while ensuring symplecticity of the map. The interpolation technique is used in an optimisation routine, to identify the lattice configuration most closely machine specified dynamical properties, including the variation of time of flight with beam energy (a key characteristic for acceleration in EMMA).

INTRODUCTION

In EMMA [1], a highly compact doublet cell is achieved using short quadrupole magnets. A large aperture requirement then leads to potentially significant fringe fields. Accurate simulations of the beam dynamics in EMMA require a dense description of the magnetic field, and numerous integration steps. Solving Maxwell's equations in an EMMA cell (by a Finite Element code, OPERA [2]) we have generated a 3D magnetic field map that can be used for numerical tracking in EMMA [3]. In most cases, numerical tracking routines are fast and reliable. However, an alternative approach based on dynamical maps could provide some benefits, particularly where speed is important; for example, when tracking a large number of particles through many cells. Dynamical maps also provide the possibility of reading significant quantities (such as tunes and chromaticities) directly from the map, giving an insight into the dynamics that is not provided directly by purely numerical methods.

To generate a dynamical map, one propagates a variable through the cell as a function instead of a numerical value. The magnetic field must be expressed in analytical form: an appropriate form can be obtained from a numerical field map by fitting an appropriate three-dimensional mode expansion [4]. Then, we use a symplectic integrator [6] implemented in the differential algebra (DA) code COSY [7], to propagate a vector of six power series (one series for each of the six dynamical variables) through the field.

* Work supported by the Engineering and Physical Sciences Research Council, UK.

† yoel.giboudot@stfc.ac.uk

DYNAMICAL MAPS CONSTRUCTION

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 of 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 [3], 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 dynamical 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 (tof), is to carry out numerical integration of the equations of motion for each new lattice configuration. However, when using dynamical maps, we can consider constructing a grid of "reference" maps, corresponding to a set of lattice configurations; we can then obtain the dynamical 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 be obtained immediately from the map, without the need for any tracking.

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 a 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 decreases 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.

Using COSY, it is possible to construct a mixed variable generating function (MVGF) that reproduces any dynamical map symplectic to order n from the map expressed in the form of a Taylor series [7]. A MVGF is not in explicit form (i.e. it cannot be used directly for particle tracking); however a map represented in the form of a 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.

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

COMPARISON WITH EXPERIMENTS

Experiments 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 behaviour as would be achieved by increasing the beam energy by 10%.

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 9.51 mm and 36.05 mm (with respect to nominal zero positions), respectively. This lattice, which we refer to as the E_1 lattice, is then specified by the notation (105.6, 100, 9.51, 36.05). Note that the baseline lattice from Berg [1] is (100, 100, 7.51, 34.05); hence in the E_1 lattice magnets are moved 2 mm outwards.

A dynamical map for the E_1 lattice was constructed using MVGF interpolation. Fig. 1 shows the comparison between measurements and simulations of the tof variation with energy for this lattice configuration. We observe that

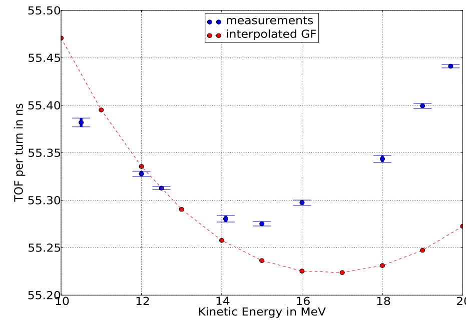

Figure 1: Comparison between measurements and simulations of the tof variation with energy for the baseline E_1 lattice configuration.

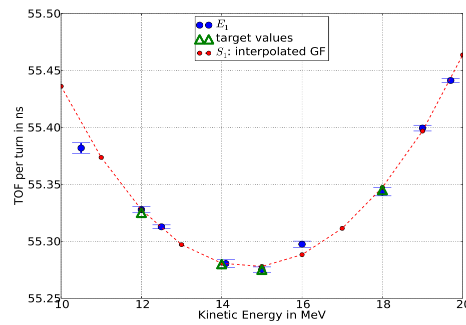

Figure 2: From the tof measurements (blue stars with error bars) four constraints (green triangles) are defined for the optimisation routine used to find a model lattice configuration with the same dynamical behaviour as observed in the machine. The tof for energies between 10 MeV and 20 MeV is then computed for the lattice found by the optimisation routine (red dashed line).

the measured lattice has a minimum tof of 55.275 ns at 15 MeV ("equivalent energy") whereas the simulated lattice has a minimum tof 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 optimisation routine was used to look for the model lattice configuration that most closely matched the measured tof curve. For each lattice configuration of the grid, we use the dynamical map to compute the tof on the closed orbit from 10 MeV to 18 MeV. Once the configuration parameters for the lattice best fitting the tof curve had been found, the dynamical map was constructed by interpolation of the MVGF.

Fig. 2 shows a comparison between measurements and simulations from the lattice configuration found by the optimisation routine. The fitted lattice found by the optimisation routine, which we refer to as S_1, has parameters (in the

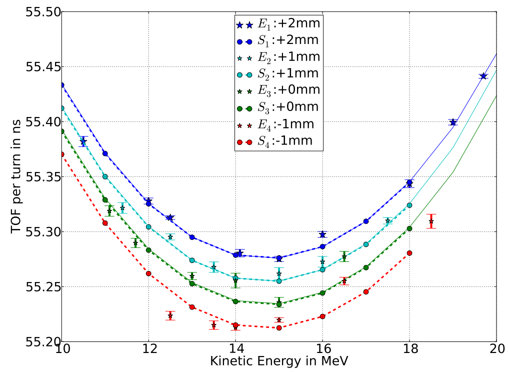

Figure 3: Comparing experimental measurements with simulation results, we observe that for lattices 1, 2 and 3, the agreement in the tof is mostly within 20 ps.

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 MeV and 20 MeV were computed for this lattice (red dashed line). It shows good agreement with measurements with a maximum discrepancy of about 20 ps at 10.5 MeV. All the other measurements agree within less than 10 ps.

When varying slightly the target values in the optimisation routine, it was possible to find lattices different from S_1 that still match relatively well the tof variation with energy. We note that the tof mainly depends on the magnet position, and has a weaker dependence on the magnet strengths. The magnet strengths can then be adjusted to agree with the tune measurements, while still matching the tof constraints. In other words a new constraint on the value of (for example) the vertical tune at high energy could be implemented in the routine. The new lattice found would correspond better to the experimental lattice. This is a possible topic for a further study.

The simulation lattice S_1 that matches the tof measurements has significantly different configuration parameters compared to the experimental lattice E_1. Although simulations and measurements do not agree in 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 tof for various energies for three transverse positions of both magnets: both moved 1 mm outward (lattice E_2, +1 mm offset), both in nominal position (lattice E_3, 0 mm offset) and 1 mm inward (lattice E_4, -1 mm offset). The magnet strengths are kept constant; for instance, the lattice E_2 is (105.6, 100, 8.51, 35.048). The tof measurements for these lattices are represented by stars with error bars in Fig. 3. We then apply the corresponding moves of the magnets 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 position of the magnets, the experimental and simulated results are plotted in the same colour.

For these simulated lattices, the tof at different energies is calculated in two different ways. The dots on the dashed line are values interpolated directly from the tof grid, hence limited to 18 MeV. The continuous line is obtained by interpolating the generating function for the corresponding lattice from the dynamical map grid. The lattice S_4 is found to be outside the grid ($X_d < 32$) and while the tof could be extrapolated, the generating function could not be derived. For the three other configurations, the generating functions agree extremely well with the interpolated tof.

We observe that for lattices 1, 2 and 3, 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). The discrepancy for the fourth lattice is larger and can be explained by the fact the configuration is outside the grid, hence the simulated data were extrapolated. However there is agreement on the minimum value of 55.22 ns between 14 MeV and 15 MeV; this value is important to determine the optimal rf frequency of the cavity to achieve acceleration.

SUMMARY AND CONCLUSIONS

The use of dynamical maps enables 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 dynamical map for any desired lattice configuration in EMMA can be obtained by interpolation on a grid of reference lattices. Basing the interpolation on mixed variable generating functions ensures the symplecticity of the interpolated map, and also appears to improve the accuracy of the map.

There appear to be significant differences between the lattice configuration used in the machine, and the lattice configuration in the model that most closely reproduces the experimental data. The reasons for these differences are not yet fully understood. However, in the cases we have looked at, we have been able to used the “fitted” model to predict the effect of changes in lattice configuration on the beam behaviour. This is an important capability in optimising machine performance, for example in adjusting the time of flight curve for maximising acceleration.

REFERENCES

[1] J.S. Berg, Nucl. Instr. Meth. Phy. Res. A 596 (2008), 276-284.

[2] Vector Field OPERA 3D. http://www.vectorfield.com/.

[3] Y. Giboudot, F. Méot, Proc. PAC’09, Vancouver, BC, Canada (2009).

[4] Y. Giboudot, R. Nilavalan, R. Edgecock, A. Wolski, Proc. IPAC’10, Kyoto, Japan (2010).

[5] S. Tygier, D. Kelliher, http://sourceforge.net/projects/pyzgoubi/

[6] Y. Wu, E. Forest, and D. Robin, Phys. Rev. E 68, 046502 (2003).

[7] COSY Infinity. http://bt.pa.msu.edu/index_cosy.htm.