Fast Symplectic Mapping and Quasi-invariants for the Large Hadron Collider*

R. L. Warnock and J. S. Berg,
Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309
É. Forest†,
Lawrence Berkeley Laboratory, University of California, Berkeley, CA 94720

Abstract

Beginning with a tracking code for the LHC, we construct the canonical generator of the full-turn map in polar coordinates. For very fast mapping we adopt a model in which the momentum is modulated sinusoidally with a period of 130 turns (very close to the synchrotron period). We achieve symplectic mapping of 10^7 turns in 3.6 hours on a workstation. Quasi-invariant tori are constructed on the Poincaré section corresponding to multiples of the synchrotron period. The possible use of quasi-invariants in deriving long-term bounds on the motion is discussed.

I. Introduction

In [1], we showed how to construct the mixed-variable generating function for the full-turn map, using only single-turn data from a symplectic tracking code. The generator is represented as a Fourier series in angle variables, the Fourier coefficients being B-spline functions of action variables. The symplectic map induced by this generator gives a good representation of the dynamics defined by the tracking code (according to physical criteria to be stated presently), even with moderate numbers of Fourier modes and spline knots. There are two special features of the B-spline–Fourier basis that promote fast map iterations: (i) the B-spline basis functions have "limited support," which is to say that only a few of the functions are non-zero at a particular point, and (ii) among all Fourier amplitudes with mode numbers less than some cutoff, a great many are found to be negligible.

A method to set long-term bounds on nonlinear motion was proposed in [2]. The idea is to make a canonical transformation to new action-angle variables $(\boldsymbol{J}, \boldsymbol{\Psi})$, such that the action $\boldsymbol{J}$ is nearly invariant, and then examine the residual variation of $\boldsymbol{J}$. In [2] the method was illustrated only in a simple example of transverse motion.

In this paper we construct maps for a realistic injection lattice of the LHC. The maps are sufficiently fast so that one can economically follow single orbits for 10^7 turns, and also construct quasi-invariant surfaces with account of synchrotron oscillations. We make the first steps toward derivation of long-term bounds, but find that the method of [2] must be elaborated if one is to find good bounds at large amplitudes.

*Work supported by DOE contracts DE-AC03-76SF00515 and DE-AC03-76SF00098

†Current address: National Laboratory for High Energy Physics, KEK, Tsukuba, Japan

II. Poincaré section at the synchrotron period

Long-term behavior of transverse coordinates is strongly affected by momentum oscillations, but the synchrotron motion itself remains roughly harmonic. For a first view of the full six-dimensional system, it is then reasonable to modulate the momentum externally, and ignore the coordinate conjugate to momentum (time-of-flight). Experience with tracking shows that such a model gives results rather similar, if not identical, to those of the full six-dimensional treatment. The synchrotron period of the LHC injection lattice is nearly 130 turns. We suppose it to be exactly 130, and take the momentum deviation at the mth turn to be

$$\delta = \frac{p - p_0}{p_0} = \delta_0 \sin\frac{2\pi m}{130}, \qquad \delta_0 = 5 \cdot 10^{-4}. \tag{1}$$

The momentum change is localized at a single r.f. cavity. For the rest of the ring we have a 4-dimensional map for fixed δ, which is represented in terms of coordinates centered at its $\delta-$dependent fixed point [1]. To save computing time in map iteration, we store the coefficients that determine the 4-d generating function, for each of the 130 values of δ.

The Hamiltonian is periodic in s with period $130C$, where C is the circumference of the reference orbit. The surface $s = 0 \pmod{130C}$ is then a Poincaré section on which there exist two-dimensional invariant surfaces and resonances that can be studied in the usual way. In comparison to the situation for the single-turn map at fixed δ, we find many more low-order resonances. This is not surprising, since even without modulation of δ the 130-th power of the map will have many more resonances than the first power. The resonance condition for the N-th power, $N\boldsymbol{m}\cdot\boldsymbol{\nu} = P$, has more solutions than that for the first power, $\boldsymbol{m}\cdot\boldsymbol{\nu} = p$, where P, p and the components of $\boldsymbol{m}$ are any integers.

III. Construction and Validation of Full-turn Map for the LHC

In contrast to the approach based on Taylor series, we do not look for a map to be valid over all of the relevant phase space. Rather, we concentrate our approximative power in a small region of action space, over which the map has relatively little variation in action. We can then get high accuracy from a small set of spline basis functions, and that allows fast iteration of the map. For a global study of stability we string together several overlapping regions, and make a map for each region.

To set the scale, we first run the tracking code to determine the short term (2000 turn) dynamic aperture. In the plane of our action variables (I_1, I_2), described in units of $10^{-7}m$, this aperture roughly follows the straight line from $(6,0)$ to $(0,9)$. To illustrate map construction, we discuss a map that is valid for initial actions $I_1(0), I_2(0)$ (with initial angles being zero) in the region R such that $1 \leq I_1(0), I_2(0) \leq 1.5$. This region is located at about one half of the short-term aperture in the plane of displacements x_1, x_2. The generator is determined in a larger region R_1, to allow for smear of orbits (as determined by short-term tracking) and an extra "apron" to account for possible long-term drift. The map iteration is programmed to stop if the orbit leaves R_1.

The map in question has 10 Fourier modes for each angle, 10 spline interpolation points for each action, and 6 for δ. The splines are cubic polynomials locally. The construction of the generator requires 264600 turns of tracking, with 68ms per turn (thus 5 hours) on an IBM RS6000 Model 590 workstation. The resulting implicit map (made explicit by Newton's method) can be iterated in 1.2ms on the same machine, giving 10^7 turns in 3.6 hours. The map agrees with the tracking code to about 1 part in 10^4 at one turn. The accuracy of agreement can be increased essentially at will by increasing the number of Fourier modes and spline points, or the order of the splines. The time for iteration does not increase with the number of spline points if the spline order is fixed (thanks to the limited support of B-splines), but the construction time is proportional to that number.

Rather than trying for higher accuracy, we consider it most interesting to work with a map of modest accuracy (hence short iteration time) and try to show that it gives essentially the same physical picture as the underlying tracking code. We do that by comparing resonances and quasi-invariant surfaces of the map and the tracking code. An easy way to find resonances (on our Poincaré surface at the synchrotron period) is to look for orbits confined to narrow bands in the plane of angles Φ_1, Φ_2; see [2]. In the case of relatively broad resonances, of which there are great many at moderate amplitudes, we always find that an initial condition giving a resonance of the map also gives the same one in tracking. In trying the same test for narrow, high-order resonances, we found a 62-nd order resonance of the map at $(I_1(0), I_2(0)) = (1.1, 1.1)$. This did not appear in tracking from the same initial condition, but another 59-th order resonance did appear. Readjusting slightly the initial condition of the map trajectory, we found the 59-th order in the map, at $(I_1(0), I_2(0)) = (1.09983, 1.1)$. This orbit of the map is plotted for 10000 synchrotron periods in Fig. 1. The corresponding orbit of tracking for 5000 synchrotron periods agrees very well on visual inspection; (quantitative comparison is difficult, since the points fill the "curves" differently in the two cases).

To compare quasi-invariant tori of the map and tracking, again on the Poincaré surface, we constructed a torus by the method of [2] in which a nonresonant orbit is fitted to a Fourier series in angle variables. Taking 20 modes for each angle, a few quasi-invariant tori of the map were computed. It typically took about 7000 synchrotron periods of mapping to compute a torus, requiring about 20 minutes on the Model 590. We check invariance under the map or under tracking by starting orbits at many randomly chosen points on the torus, and see how close the orbit is to the torus after one synchrotron period. With 50 randomly chosen points, a typical surface was invariant to one part in 10^5 under the map from which it was constructed, and invariant to

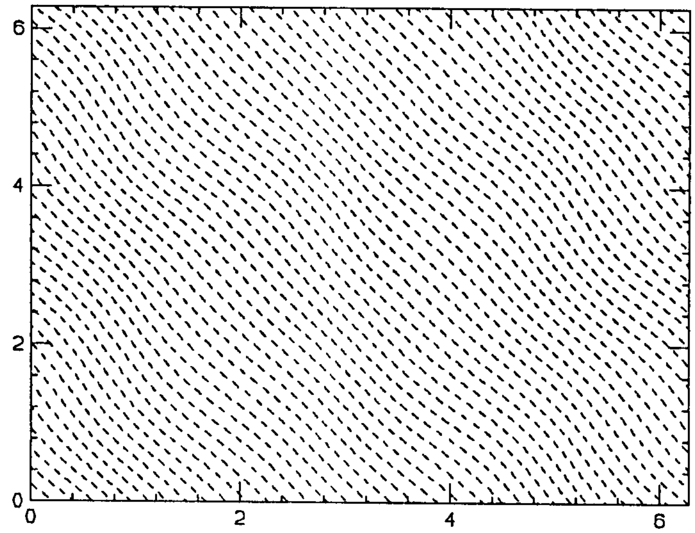

Figure. 1. Φ_2 vs. Φ_1 on a 59th order resonance of the map, $33\nu_1 \pm 26\nu_2 = p$, for 10000 synchrotron periods. Initial conditions: $I_1(0)$, $I_2(0) = (1.09983, 1.1) \cdot 10^{-7}$ m.

one part in 10^4 under the tracking code.

These and other tests convince us that the Hamiltonian system represented by symplectic maps of modest one-turn accuracy (1 part in 10^4) represent a physical system very similar to that of the underlying tracking code, at least at amplitudes not too close to the dynamic aperture. At very large amplitudes it is not easy to validate the map by the above arguments, since one finds large-scale chaotic behavior rather than clean resonances and quasi-invariants.

IV. A String of Large-amplitude Maps and Long-term Mapping

In one over-night run, using a small fraction of available CPU time on a "farm" of workstations at SLAC, we produced a string of five maps in partially overlapping rectangular regions. Numbers of modes and spline knots were the same as in the example above. This gives a continuous strip of allowable initial conditions (with $\Phi_1 = \Phi_2 = 0$), between two lines running parallel to the short-term aperture. The outer border of the strip is at 70% of the aperture. The outer corners of the rectangles go beyond the strip, and allow orbits that go within at least 85% of the aperture. At the time of writing we have done a few runs of 10^7 turns using these maps. Fig. 2 shows a plot of I_1 at every eighth synchrotron period in such a run. The vertical frame size of the graph indicates the domain of the splines in I_1. The domain of the maps does not include the coordinate axis, $I_1 = 0$, $I_2 = 0$. The map construction fails in a small neighborhood of each axis, since the polar coordinate system becomes inappropriate.

V. A First Try at Long-term Bounds

Here we are concerned solely with the dynamics defined by the map. For a first attempt at long-term bounds we work at smaller amplitudes, in a region with $0.4 < I_1(0), I_2(0) < 0.6$ (in units of $10^{-7}m$). Following the method of [2], we construct

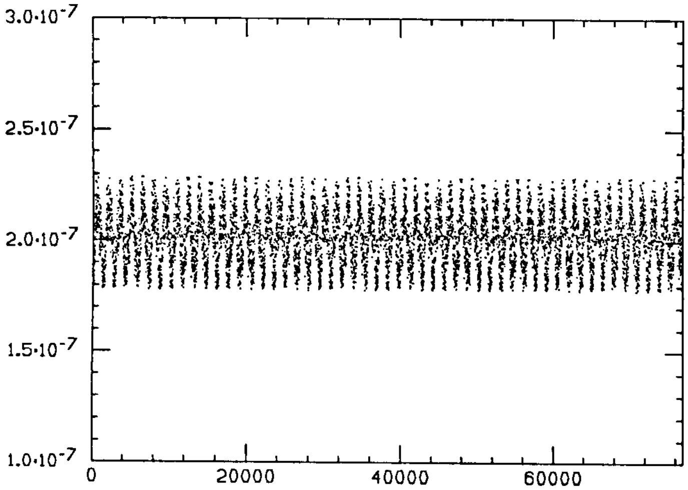

Figure. 2. I_1 vs. number of synchrotron periods, for 10^7 turns. Initial conditions: $I_1(0)$, $I_2(0) = (1.87, 0.937) \cdot 10^{-7}$ m.

a set of 9 tori on the Poincaré section, for points close to a 3×3 rectangular grid in $I_1(0), I_2(0)$. These tori have twenty Fourier modes in each variable, and are invariant to about 1 part in 10^5. We next interpolate the tori in actions, so as to define a smooth canonical transformation to new action-angle variables $\boldsymbol{J}, \boldsymbol{\Psi}$. Although $\boldsymbol{J}$ is fairly constant on and near the original tori, it remains to be seen how much it varies in the region of interpolation.

Let Ω denote the region in which $\boldsymbol{J}$ is defined. Suppose that δJ is an upper bound for the change in $|J_i|$ during m synchrotron periods, for any orbit beginning in Ω. Let Ω_o be a subregion of Ω, and let ΔJ be the minimum distance from Ω_o to the boundary of Ω. Then an orbit beginning in Ω_o cannot leave Ω in fewer than nm synchrotron periods, where $n\delta J = \Delta J$. Then we have stability (in the sense of being confined to Ω) for $N = 130m\Delta J/\delta J$ turns.

We take $m = 1000$ and try 1000 randomly chosen initial conditions in Ω to estimate δJ. We find that δJ is about 0.01, which is much larger than the variation of J for orbits starting on the original tori. If we take Ω_o to be a small box in the middle of Ω, then ΔJ is about 0.1, and we can predict stability only for $N = 1.3 \cdot 10^6$ turns, a disappointingly small number.

The reason for the large variation of $\boldsymbol{J}$ is the presence of a fairly broad resonance inside Ω. A sufficiently isolated resonance in mode $\boldsymbol{m} = (m_1, m_2)$ can be identified by plotting the change of $\boldsymbol{m} \cdot \boldsymbol{J}$ against $\boldsymbol{m} \cdot \boldsymbol{\Psi}$ at constant $K = m_1 J_2 - m_2 J_1$ [2]. As is shown for a $(6, 1)$ resonance in Fig. 3, these variables perform a pendulum-like motion. Note that this motion would be hard to see without first transforming to $\boldsymbol{J}, \boldsymbol{\Psi}$. The resonance could be quite stable, but still lead to large oscillations in $\boldsymbol{J}$. In order to make our argument for long-term bounds, it will be necessary to find quasi-invariants of resonant orbits. Preliminary work by Armando Antillón showed that a simple pendulum-type Hamiltonian $H(\boldsymbol{m} \cdot \boldsymbol{J}, \boldsymbol{m} \cdot \boldsymbol{\Psi}, K)$ could be fitted to orbit data so as to provide at least a rough quasi-invariant; K is a second quasi-invariant. We hope to use H and K in place of J_1, J_2 in our argument for long-term bounds. It may be necessary to refine the definition of these quantities, so that they are functions of all the canonical variables $\boldsymbol{J}, \boldsymbol{\Psi}$.

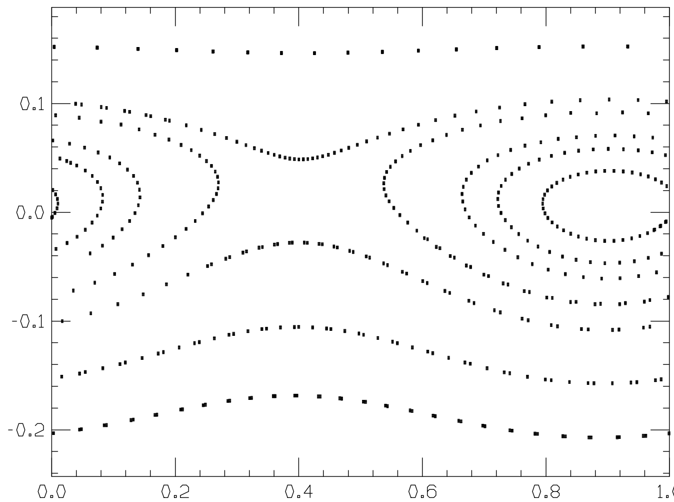

Figure. 3. Plot of $\hat{\boldsymbol{m}} \cdot d\boldsymbol{J}$ vs. $[\boldsymbol{m} \cdot \boldsymbol{\Psi} \pmod{2\pi}]/2\pi$, where $\hat{\boldsymbol{m}}$ is the unit vector in the direction of $\boldsymbol{m} = (6, 1)$, and $d\boldsymbol{J}$ is the deviation of $\boldsymbol{J}$ from a fixed "average action," $\boldsymbol{J}_0$. Each of the 9 orbits plotted is started at the same value of $K = m_2 J_1(0) - m_1 J_2(0)$. The action unit is 10^{-7} m.

VI. Acknowledgment

This work was made possible by the kind help of Frank Schmidt, who provided lattice information for the LHC and assistance in verifying our tracking code.

References

[1] J. S. Berg, R. L. Warnock, R. D. Ruth, and É. Forest. *Physical Review E*, 49:722–739, January 1994.

[2] R. L. Warnock and R. D. Ruth. *Physica D*, 56:188–215, 1992.