SLAC-PUB-96-7091
January 1996

FAST SYMPLECTIC MAPPING, QUASI-INVARIANTS, AND LONG-TERM STABILITY IN THE LHC *

ROBERT L. WARNOCK and J. SCOTT BERG

Stanford Linear Accelerator Center, Stanford University
Stanford, California, 94309

A systematic program to explore stability of orbits in hadron storage rings is based on the following steps: (a) beginning with a symplectic tracking code, construct the mixed-variable generator of the full-turn map in a Fourier-spline basis; (b) use the resulting fast mapping to follow long orbits and estimate the long-term dynamic aperture; (c) contruct quasi-invariants and examine their variation in time to set long-term bounds on the motion for any initial condition in a specified region. First results from an application of the program to the Large Hadron Collider (LHC) are reported. Maps can be constructed in a few hours and evaluated at a speed 60 times greater than that of one-turn tracking, on a workstation computer. Orbits of 10^7 turns take 3.6 hours. The value of a "stroboscopic" view of the synchro-betatron motion is emphasized. On a Poincaré section at multiples of the synchrotron period, one can study resonances and invariant surfaces in two dimensions, thereby taking advantage of techniques that have proved effective in treating pure betatron motion.

KEY WORDS: accelerators, nonlinear mechanics, symplectic maps, symplectic integrators, invariant tori, nonlinear resonances

1 THE FULL-TURN MAP AS DYNAMICAL MODEL

The most reliable simulation of single-particle motion in accelerators is obtained from element-by-element symplectic tracking, using a realistic lattice. In the case of large hadron storage rings this direct approach is limited by computation cost, since one cannot afford to follow many orbits over the desired storage time of the beam. Effectively, tracking does not provide a dynamical model for large times. Three remedies have been suggested: (i) replace the realistic lattice by a much simpler one[1] allowing faster tracking; (ii) avoid a real definition of long-term dynamics, trying instead to derive some signal of long-term behavior from behavior over shorter times (for instance, by looking at deviation of two initially close orbits[2], or examining pseudo-tunes obtained from windowed time series[3]); (iii) construct a full-turn map that approximates the dynamics of a realistic tracking code, while giving a much faster calculation of turn-by-turn evolution. Although methods (i) and (ii) might

* Work supported by Department of Energy contract DE–AC03–76SF00515.

To appear in the proceedings of the International Workshop on Single Particle Effects in Large Hadron Colliders, Montreux, Switzerland, October 15–21, 1995.

find a useful role, they are subject to irreducible uncertainties. In our opinion, the mapping method (iii) can be nearly as reliable as direct tracking, provided that an appropriate technique is applied.

The usual habit in the field of nonlinear mechanics is to favor techniques that derive from classical analytical methods such as power series and perturbation theory, even when the calculation is being done numerically. The spirit of modern numerical analysis, with its emphasis on *local* approximation and interpolation with controlled error, is rarely noticed. We argue that methods suggested by numerical analysis are appropriate both for representation of maps and for computation of invariants and canonical transformations.

Maps based on Taylor series have not been entirely successful for the study of long-term evolution. (Their use in producing normalizing transformations is a different matter). They typically do not satisfy the symplectic condition with adequate accuracy at large amplitudes. Also, the construction time for Taylor series of suitably high order is uncomfortably large in the case of big machines. The matter of symplecticity can be handled by deriving a mixed-variable generating function from the Taylor series. Yan *et al.*[4] derived the generator by formal power expansions, and applied it successfully to mapping of orbits in the SSC. Another approach that is being tried is to represent the map as a composition of a relatively small number of nonlinear "kicks" and linear symplectic transformations[5]. This form is symplectic, but it may be difficult to understand its convergence properties.

The Taylor development about the origin has doubtful efficiency as a representation of either the map or its generating function at points far from the origin, in particular for interesting orbits close to the dynamic aperture. To get an efficient local approximation for large amplitudes, one should use polar coordinates, since the radial (action) variable typically has relatively small variation along an orbit. The action-dependence of the map can be represented locally in terms of a few simple basis functions. A bonus is that the angle dependence proves to be fairly simple as well, in the sense that relatively few Fourier modes are appreciable. A disadvantage of polar coordinates is that they lead to singularities in generating functions at zero action; we must stay away from the origin in every phase plane, or else use a mixed Cartesian-polar system if one phase-plane point is near the origin.

2 CONSTRUCTION OF THE MAP FROM A TRACKING CODE

To enforce the symplectic condition, the full-turn map is defined implicitly by a mixed variable generator $G(\mathbf{I}, \mathbf{\Phi}')$, where the components of vectors $\mathbf{I}$ and $\mathbf{\Phi}'$ are action and angle variables, respectively, of the underlying (normalized) linear system. For each iteration of the map, the explicit map $T : (\mathbf{I}, \mathbf{\Phi}) \mapsto (\mathbf{I}', \mathbf{\Phi}')$ is obtained from G by Newton's method. The construction of G (from single-turn tracking data for many initial conditions) is described in detail in Ref.([6]). Although G is a nonlinear function of the tracking data, the problem of constructing G can be solved through Fourier analysis and a nonlinear change of variable.

For application to accelerators we let G represent the map M for all of the ring except the r.f. cavity. The fixed point of this map depends on the total momentum deviation $\delta = -(p - p_0)/p_0$, and therefore G is represented in terms of transverse coordinates $(I_1, I_2, \Phi_1', \Phi_2')$ referred to an origin at the δ-dependent fixed point. The total map is composed of a simple explicit map for the cavity, a linear transformation to center the coordinates at the fixed point, and M. For details see Ref.([6]).

The generator is represented as a Fourier series in $\mathbf{\Phi}'$, with coefficients given as B-spline functions of $\mathbf{I}$ and δ. The B-spline basis promotes fast evaluation of the map, since only a few of the basis functions are non-zero at any point. Further gains in speed are achieved by dropping negligible Fourier modes and by starting the Newton iteration at a point obtained from a crude explicit map.

3 A STROBOSCOPIC VIEW OF SYNCHRO-BETATRON MOTION

For a first exploration we take advantage of the fact that the betatron motion does not have a strong effect on the synchrotron motion. We represent the latter by a sinusoidal modulation of δ, and ignore the coordinate conjugate to δ (time of flight). If we approximate the synchrotron tune by a rational number, the Hamiltonian is then periodic in s with period equal to some multiple of the ring circumference C. We study the LHC in injection mode, for which the synchrotron tune is $\nu_s = 1/129.97$. Approximating by $\nu_s = 1/130$, we have a periodicity of 130 turns, and a four-dimensional Poincaré section at $s = 0 \pmod{130C}$.

If we view an orbit only as it intersects this four- dimensional surface, once per synchrotron period, then we have a fruitful way of analyzing the motion[7]. One can study resonances and two-dimensional invariant tori of the 130-turn map, and set long-term bounds by looking at two-dimensional quasi- invariant actions. To evaluate the 130-turn map we merely iterate the single-turn map 130 times. To save time in iteration, we store coefficients of the generator for all 130 values of δ.

4 VALIDATION OF THE MAP, AND RESULTS ON SPEED OF MAPPING

Our approximation of the map has good convergence characteristics, so that one can achieve high accuracy by increasing the number of Fourier modes and spline interpolation points. For high speed of map iteration it is more interesting, however, to see what can be accomplished with a map of modest accuracy. Our map will always represent a Hamiltonian system (modulo round-off error), since symplecticity is not compromised. This system may closely resemble the Hamiltonian system of the tracking code, even if orbits of the two systems starting at the same point do not agree closely at large time. If the "skeleton of phase space," defined by the geometric structures of resonances and invariant tori, is the same down to some fine scale, it is reasonable to assume that the map provides a good model of tracking.

Here and in Ref.([9]), we give some results on maps which include Fourier modes up to $m = 8$, and 10 action interpolation points for cubic splines, in either degree of freedom; there are 6 interpolation points in δ. The maps are defined in rectangles in action space. Maps for a sequence of overlapping rectangles cover a strip in action space running parallel to the short-term (2000 turn) dynamic aperture, at about 60% of that aperture. They agree with tracking to about one part in 10^4 at one turn, and can be iterated for 10^7 turns in 3.6 hours on an IBM RS6000-590 workstation.

One can construct an invariant surface (in the Poincaré section at the synchrotron period) in about twenty minutes on average. We find that a surface constructed from the map, invariant to one part in 10^5 under the map, is invariant to one part in 10^4 under tracking. Although phase error may build up in the comparison of mapping and tracking orbits, the orbits nevertheless adhere closely to the same torus. Similarly, resonances of mapping show close similarity to those of tracking. Resonances are easily identified by plotting one angle variable against the other, $\pmod{2\pi}$. Such a $\Phi_1 - \Phi_2$ plot would show straight lines for resonances of an integrable system. In our case the lines become wavy, and may thicken to narrow bands if the resonance is broad. The number of intersections of lines (bands) with the axes determine the mode numbers of the resonance (modulo a relative sign). Finding a low order resonance in tracking by this method, we find the same pattern in the angle plots from mapping. High-order resonances provide a more demanding test. A 59-th order resonance found in tracking did not appear in mapping initially, but after a slight adjustment of initial condition it was found. At high amplitudes where motion is chaotic it is not possible to validate the map by such tests, but one can still look for resonances imbedded in chaos (not an infrequent phenomenon) or try to compare borders of chaotic regions.

More tests should be carried out, but our work to date suggests that the uncertainty in using these maps in place of tracking is very much less than the primary uncertainty in the magnetic fields that define the Hamiltonian.

5 CONSTRUCTION OF QUASI-INVARIANTS FROM ORBIT DATA

Let $\mathcal{P}_n = \{(\mathbf{I}^{(i)}, \mathbf{\Phi}^{(i)}), i = 1, 2, \cdots, n\}$ be a sequence of points in the Poincaré section $\mathcal{S}$ at the synchrotron period, all points lying on a single orbit. For a large set of initial conditions, the orbit at $\mathcal{S}$ will lie on an invariant torus, which can be represented as a Fourier series for the action variable,

$$\mathbf{I} = \sum_{\mathbf{m}} \mathbf{I}_{\mathbf{m}} e^{i\mathbf{m}\cdot\mathbf{\Phi}} \ . \tag{1}$$

Different tori are labeled by the amplitude of the zero mode, $\mathbf{J} = \langle \mathbf{I} \rangle$, otherwise known as the invariant action. The orbit is transitive on the torus, which is to say that some $\mathbf{\Phi}^{(i)}$ comes close to any point $\mathbf{\Phi}$ for sufficiently large i. On a resonant orbit there are regions of $\mathbf{\Phi}$-space not approached by any $\mathbf{\Phi}^{(i)}$, and the relation (1) between $\mathbf{I}$ and $\mathbf{\Phi}$ does not hold.

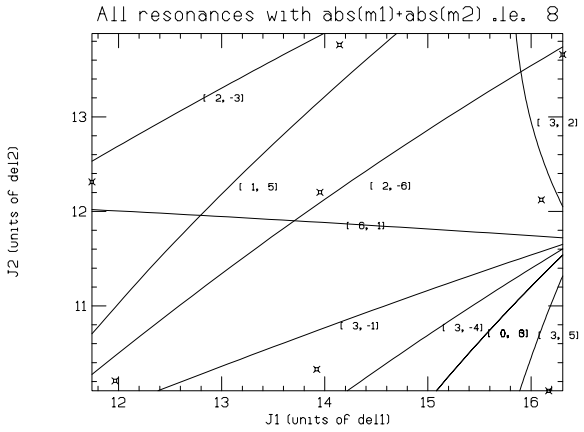

FIGURE 1: Resonance lines in $J_1 - J_2$ plane, units of 10^{-8}m. Points marked with an asterisk correspond to the original tori. One of the 9 points is outside the box, near the upper left corner.

When Eq.(1) holds, it defines the generating function $\mathcal{G}$ of a canonical transformation $(\mathbf{I}, \mathbf{\Phi}) \mapsto (\mathbf{J}, \mathbf{\Psi})$, where $\mathbf{J}$ is constant, and $\mathbf{\Psi}$ advances linearly with the time. Indeed,

$$\mathbf{I} = \mathbf{J} + \mathcal{G}_{\mathbf{\Phi}}(\mathbf{J}, \mathbf{\Phi}) = \mathbf{J} + \sum_{\mathbf{m}} i\mathbf{m} g_{\mathbf{m}}(\mathbf{J}) e^{i\mathbf{m}\cdot\mathbf{\Phi}} \ . \tag{2}$$

For an orbit on an invariant torus, the Fourier coefficients $\mathbf{I}_{\mathbf{m}}$ (hence $\mathbf{J}$ and $g_{\mathbf{m}}$) may be determined from the sequence $\mathcal{P}_n$, with arbitrary accuracy for large n. Since the $\mathbf{\Phi}_n$ are scattered unpredictably, the determination is not a standard problem of numerical Fourier analysis, which requires uniformly distributed data. A stable method for scattered data[8] is based on the idea of using the values of $\mathbf{I}$ on a uniform mesh as the unknowns; those values are related to the Fourier amplitudes by a discrete Fourier transform. We take n sufficiently large so that there is at least one $\mathbf{I}^{(i)}$ in each cell of the mesh, then throw away all except one value per cell. The resulting large system of linear equations for $\mathbf{I}$ on the mesh is soluble by iteration. If the orbit is resonant and the mesh is fine, there will always be empty cells; as it should, the algorithm then fails to produce a torus.

We use this method to determine the $g_{\mathbf{m}}(\mathbf{J})$ on a number of tori, then interpolate smoothly in $\mathbf{J}$ to define a canonical transformation globally. This interpolation bridges over resonances. Correspondingly, interpolated values of $\mathbf{J}$ are not always as constant in time as the $\mathbf{J}$ on the original tori. Tunes (of the map for a synchrotron period) are obtained from the advance of $\mathbf{\Psi}$ per period. For the interpolated tori one has pseudo-tunes, interpolating the genuine tunes of the original tori. The interpolated tune is normally an invertible function of $\mathbf{J}$, from which one can find resonance lines in the $\mathbf{J}$ plane. An example is shown in Figure 1. This diagram of course depends somewhat on the chosen interpolation; i.e., on the definition of $\mathbf{J}$.

This method of finding invariants and canonical transformations is quite robust and automatic. Since it succeeds at large amplitudes it is preferable to perturbation

theory in applications such as those of the following sections. Perturbation theory can diverge at amplitudes beyond the first island chain, even if perfectly good tori exist in that region. Another advantage, important in the present context, is that one can find invariant tori of a high power of the map without having an explicit formula for that power. Our method deals with the small divisor problem simply by filtering out resonant orbits. When a resonant orbit is detected through persistently empty cells in $\mathbf{\Phi}$ space, our program automatically tries a new orbit. Since we typically use at least a 41×41 mesh in the $\Phi_1 - \Phi_2$ plane, to get 20 Fourier modes in each direction, resonances of rather high order are filtered out. Even so, the number of rejected orbits proves to be manageable.

6 LONG-TERM BOUNDS ON THE MOTION

Since $\mathbf{J}$ is much more constant in time than the original action coordinate $\mathbf{I}$, its residual variation provides a sensitive indicator of long-term drift of an orbit.[10] Suppose that δJ is an upper bound for the change in $|\mathbf{J}|$ during M synchrotron periods, for any initial condition in a region Ω of phase space. If Ω_o is a subregion of Ω, and the minimum distance from the boundary of Ω_o to the boundary of Ω is ΔJ, then an orbit originating in Ω_o cannot leave Ω in fewer than pM periods, where $p\delta J = \Delta J$. Thus, orbits starting in Ω_o are stable, in the sense of being confined to Ω, over $N = Mn_s\Delta J/\delta J$ turns, where n_s is the number of turns in a synchrotron period. We would like to have an N comparable to the desired number of turns for beam storage, with ΔJ representing an acceptable action excursion in operation of the machine. One difficulty in applying this argument is to get a reliable value for δJ with suitably large M. One can get a plausible value from random sampling of initial conditions in Ω.[10] A rigorous determination is possible in principle, through interval arithmetic.[11]

7 PENDULUM MOTION, QUASI-INVARIANTS, AND LONG-TERM BOUNDS NEAR A STRONG RESONANCE

For applications in regions Ω near the short-term dynamic aperture, one finds that δ may be too large to get a large enough value of N. (Here we must restrict M for feasible computational expense; $M = 500$ to 10000 is typical). The δJ may be suitably small on the original tori determined from non-resonant orbits, while being too large in interpolated regions. This may come about because the interpolation bridges over a strong resonance. For instance, δJ near the $(6, 1)$ resonance of Figure 1 is about 100 times larger than δJ on the original tori.

It is remarkable that the motion near a strong resonance closely resembles pendulum motion when viewed in appropriate coordinates. If we plot $\mathbf{m} \cdot \mathbf{J}$ versus $\mathbf{m} \cdot \mathbf{\Psi} \pmod{2\pi}$ on the section $\mathcal{S}$ we see the familiar libration curves (angle restricted to a subinterval of $[0, 2\pi]$) and rotation curves (angle ranging over $[0, 2\pi]$).

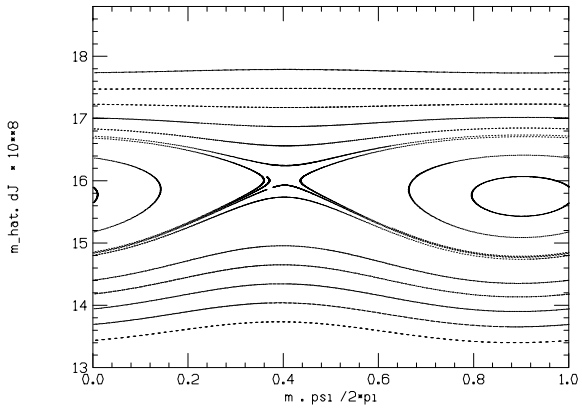

FIGURE 2: Pendulum motion in plot of $\hat{\mathbf{m}} \cdot \mathbf{J}$ vs. $\mathbf{m} \cdot \mathbf{\Psi} \pmod{2\pi}$, units of 10^{-8}m

This is shown in Figure 2 for the resonance with $\mathbf{m} = (6, 1)$ that appeared in Figure 1. Correspondingly, $\mathbf{m} \times \mathbf{J}$, the component of $\mathbf{J}$ perpendicular to $\mathbf{m}$, is nearly constant. In other words, the motion resembles that of an isolated-resonance Hamiltonian, which has the form

$$H = \sum_{n \in Z} h_n(\mathbf{m} \cdot \mathbf{J}, \mathbf{m} \times \mathbf{J}) \exp(in\mathbf{m} \cdot \mathbf{\Psi}) \ . \tag{3}$$

This simple behavior cannot be seen in the original variables $(\mathbf{I}, \mathbf{\Phi})$, even if we plot $\mathbf{m} \cdot \mathbf{I}$ against $\mathbf{m} \cdot \mathbf{\Phi}$. It is necessary to remove what might be called the "principal smear" or "non-resonant background" by the canonical transformation $(\mathbf{I}, \mathbf{\Phi}) \mapsto (\mathbf{J}, \mathbf{\Psi})$.

A closer look naturally reveals that we do not have exact pendulum motion. There is some scatter in the value of $\mathbf{m} \times \mathbf{J}$, and a similar amount of scatter in the points which seem at first to define rotation or libration curves. Moreover, an orbit near the boundary of the rotational region may first look like a separatrix and then look like libration near a separatrix. Thus, we have a mechanism for fast transport; an orbit can cross an island if it is sufficiently close to the island. One of the orbits in Figure 2 is of this sort (its librational part is not plotted for the full librational range; the plot ends near the turning point).

Although a relatively small set of orbits display fast transport, orbits well inside or well outside the island appear to be stable over a long time. For a "worst case" estimate, we can estimate stability time for orbits in the rotation regions above and below the island, and merely assign zero time of transport across the island for orbits near the island. Any orbit that drifted upward from the upper pseudo-separatrix region or downward from the lower would fall into a rotational motion with long survival time.

Among all the low-order resonances shown in Figure 1, only the (6,1) is strong enough to be easily found in orbit data. Moreover, the (6,1) makes itself felt over

the entire region of the (J_1, J_2) plane indicated in Figure 1, in that a plot of $\mathbf{m} \cdot \mathbf{J}$ versus $\mathbf{m} \cdot \mathbf{\Psi}$ with $\mathbf{m} = (6, 1)$ shows a rotation or libration curve for a large set of initial $\mathbf{J}$ chosen randomly in the region. Note that this is quite a big region, in which the components of $\mathbf{J}$ vary by almost 40% .

For a formal argument on long-term bounds we take the quasi-invariants near a strong resonance to be $K_1 = \hat{\mathbf{m}} \times \mathbf{J}$ and $K_2 = \langle \hat{\mathbf{m}} \cdot \mathbf{J} \rangle$, where K_2 is the zero mode amplitude in a Fourier series fitted to $\hat{\mathbf{m}} \cdot \mathbf{J}$ (on rotation curves only) as a function of $\mathbf{m} \cdot \mathbf{\Psi}$. Here $\hat{\mathbf{m}}$ is the unit vector in the direction of $\mathbf{m}$. In the region $\Omega = \{\mathbf{J}, \mathbf{\Phi} | 12.5 < J_1 < 15.5 \ , \ 10.5 < J_2 < 13.5 \ , \ \Phi_i \in [0, 2\pi]\}$, with actions in units of 10^{-8}m, we estimate δK_1 to be less than $4 \cdot 10^{-11}$m for 500 synchrotron periods. We have estimated δK_2 less carefully, but it appears to have a similar magnitude. These values would indicate that an orbit starting near the middle of Ω would not leave Ω in fewer than 10^7 turns or so. Since the region in question is at about 60% of the short term aperture, we conclude that useful long-term bounds on the motion are feasible at fairly large amplitudes, even in the neighborhood of a strong resonance. Better quasi-invariants (hence bounds for longer times) might be obtained by allowing dependence on the angle $\mathbf{m} \times \mathbf{\Psi}$, for instance by defining K_2 as the zero mode in a two-dimensional Fourier series in the two angles. We intend to report more careful and comprehensive calculations in a later communication.

ACKNOWLEDGEMENTS

We thank Frank Schmidt and Étienne Forest for help in providing a tracking code for the LHC.

REFERENCES

1. D. Ritson, p.304 in *Nonlinear Problems in Future Particle Accelerators*, W. Scandale and G. Turchetti, eds. (World Scientific, 1991).
2. F. Galluccio and F. Schmidt, *3rd European Part. Accel. Conf.*, CERN LHC Note 184, 1992.
3. W. Scandale, private communication.
4. Y. Yan, P. Channell, and M. Syphers, SSCL-Preprint-157, 1992.
5. J. Irwin, SSC-228, 1989. D. T. Abell and A. Dragt, these Proceedings.
6. J. S. Berg, R. L. Warnock, R. D. Ruth, and É. Forest, Phys. Rev. E **49**, 722 (1994).
7. We understand that S. Peggs and L. Evans independently suggested this stroboscopic picture. As far as we know, they did not treat resonances and tori on the Poincaré section.
8. R. L. Warnock, Phys. Rev. Lett. **66**, 1803 (1991).
9. R. L. Warnock and J. S. Berg, *1995 IEEE Part. Accel. Conf.*, SLAC-PUB-95-6840.
10. R. L. Warnock and R. D. Ruth, Physica D **56**, 188 (1992).
11. G. H. Hoffstätter, DESY 94-242, 1994; M. Berz and G. H. Hoffstätter, Interval Computations, 1994.