A. Wittig, M. Berz

HIGH PERIOD FIXED POINTS, STABLE AND UNSTABLE MANIFOLDS, AND CHAOS IN ACCELERATOR TRANSFER MAPS

Michigan State University, East Lansing, Michigan, 48824, USA.

In this paper we present an algorithm for a verified global fixed point finder. More specifically, a method is described to automatically identify and classify regions of the search space which are guaranteed to either contain none, precisely one, or one or more fixed points, as well as regions that may or may not contain fixed points. The fixed point finder is implemented with Taylor Models in COSY INFINITY, allowing for very efficient identification of fixed points even in numerically complicated systems with high dependency and strong cancellation.

We then apply the fixed point finder to find higher order periodic points in a transfer map taken from the Tevatron accelerator. The results are compared to predictions made from tune shifts computed using normal form theory. A high order approximation to the stable and unstable manifolds of a set of hyperbolic periodic points is computed and shown.

Keywords: Taylor model, fixed points, chaos, manifolds.

Introduction. The study of fixed points and periodic points of maps, such as transfer maps for storage rings and synchrotrons, can reveal much information about the behavior of such a dynamical system. This is particularly true for beam transfer maps, as it is often the case that undesirable chaotic behavior in an accelerator is observed near periodic points. So automated detection of such periodic points can be of great help in determining dynamic apertures more rigorously.

In order to do so, it is of course first necessary to find the fixed and periodic points of a given map. For all but the very simplest maps, it is not possible to derive a closed form analytic solution for all fixed points or periodic points of a given period. Leveraging the power of modern computers, the use of numerical methods can be of great help in determining fixed and periodic points. Simple numerical techniques, such as Newton methods, have been around for a long time and are a well known tool for such tasks. When applied to a map using floating point numbers readily available in computer hardware, they can quickly identify isolated, approximate fixed points.

However, these techniques suffer from two major drawbacks. Given a starting point, they converge to a single fixed or periodic point approximation, but there is no information obtained about possible other fixed or periodic points in a given area. Secondly, floating point numbers are not mathematically rigorous due to round-off errors caused by the finite precision of floating point numbers. Therefore, even if a floating point candidate of a possible

fixed or periodic point is identified, it is not guaranteed that the underlying map really exhibits the same fixed or periodic point.

In the following, we present an algorithm for a verified, global fixed and periodic point finder. That is, we will identify regions within a prescribed search space where we can say with certainty that there are no fixed or periodic points. For the remaining regions, we will perform rigorous tests to determine if they include exactly one, or possibly several, fixed or periodic points. As a result, all fixed points and periodic points up to a pre-specified period will be found with certainty.

To perform rigorous computations, we use the technique of Taylor Models [1] as implemented in the COSY INFINITY [2] rigorous computation package. It will be shown that Taylor Models are particularly well suited for this type of problem, and yield superior results to other rigorous computational techniques such as interval arithmetic.

We then proceed to illustrate the algorithm by applying it to a real world transfer map of the Tevatron accelerator at Fermilab. This map is represented as a polynomial of order twenty, presenting particular numerical challenges to overcome.

Fixed Points and Their Identification. Let $K \in \mathbb{R}^m$ be a compact set and $M : K \to \mathbb{R}^n$, $n \in \mathbb{N}$, a sufficiently often continuously differentiable map. Ideally, we are interested in finding all periodic points of given order $p > 0$ in a compact subset. That is, we want to find all $\mathbf{x} \in K$ such that

$$M^p(\mathbf{x}) = \mathbf{x}. \tag{1}$$

Due to the finite precision of numerical computations, it is in general not possible to find all such fixed points precisely. Therefore we will slightly reformulate the problem. Instead of computing all possible solutions $\mathbf{x}_i \in K$ exactly, we will determine small compact subsets $X_i \subset K$ such that for some $\mathbf{x}_i \in X_i$ Equation (1) holds. It will be our goal to make the X_i sufficiently small in measure to get a good enclosure of the real periodic point.

We then further classify each subset X_i by whether it contains exactly one, or one or more periodic points. As is typical for verified computation there is also the possibility of remaining subsets about which no statement can be made at all. That is, there is possibly a third category of sets for which we do not know if they contain any periodic points at all.

Furthermore, since the case $p > 1$ coincides with $\mathbf{x}$ being a fixed point of the map M^p, without loss of generality we will only consider fixed points in the following.

Classification. We begin by describing the mathematical foundation for classifying a single subset $C \subset K$ according to the above criteria. To that end, we recall two important theorems of Analysis, the Schauder fixed point theorem [3] and the Banach fixed point theorem [4]:

Theorem 1 (Schauder's fixed point theorem). *Let $C \subset \mathbb{R}^n$ be non-empty, convex, and compact, and f be a continuous map from $C \to C$. Then*

$$\exists \mathbf{x}_i \in C \text{ such that } f(\mathbf{x}) = \mathbf{x}.$$

Theorem 2 (Banach's fixed point theorem). *Let C be a non-empty, complete metric space and $f : C \to C$ be a contraction. Then*

$$\exists! \mathbf{x} \in C \text{ such that } f(\mathbf{x}) = \mathbf{x}.$$

Recall that f is a contraction if $\exists k, 0 \leqslant k < 1$ such that

$$|f(\mathbf{x}) - f(\mathbf{y})| < k |\mathbf{x} - \mathbf{y}| \quad \forall \mathbf{x}, \mathbf{y} \in C.$$

The requirements of Theorem 1 are well suited for rigorous numerical computations. Computing an enclosure of the image of a suitable set C is a straightforward task for any verified numerics packages, as is bounding an checking for inclusion.

Note that Theorem 2 can be understood as an extension of Theorem 1. If the all prerequisites of Theorem 1 are met, and thus existence of a fixed point in C is already proved, uniqueness is obtained by the additional requirement that f be a contraction on C.

Verification. The verification of the contraction property, unfortunately, is not quite as straightforward as the verification of the inclusion mapping property in Theorem 1. In the following, we will provide an overview over various ways to assert the contractive property of f, and discuss their fitness for verified numerical computations.

We begin with the following definition of the operator norm of a matrix:

Definition 1 (Matrix Operator Norm). Let A be a $n \times n$ matrix. The matrix operator norm of A is given by

$$||A|| = \sup_{|\mathbf{x}|=1} |A \cdot \mathbf{x}| = \sup_{\mathbf{x}\neq 0} \frac{|A \cdot \mathbf{x}|}{|\mathbf{x}|}$$

where the vector norm $|\cdot|$ is the Euclidean norm on $\mathbb{R}^n$. It is easy to see that due to linearity both expressions are equivalent.

Lemma 1. *Let $C \subset \mathbb{R}^n$ be a non-empty, convex, and compact set and f by a continuously differentiable function on C such that $||Df(\mathbf{x})|| < k < 1 \; \forall \mathbf{x} \in C$. Then f is a contraction on C.*

Proof. Given any two distinct points $\mathbf{x}, \mathbf{y} \in C$, consider the function $G : [0,1] \to \mathbb{R}$ given by the dot product

$$G(t) = (f(\mathbf{y}) - f(\mathbf{x})) \cdot f\left((1-t)\mathbf{x} + t\mathbf{y}\right).$$

Applying the Mean Value Theorem to this function, one finds that for some $\xi \in (0,1)$

$$(f(\mathbf{y}) - f(\mathbf{x})) \cdot (f(\mathbf{y}) - f(\mathbf{x})) = G(1) - G(0) =$$
$$= \; G'(\xi) = (f(\mathbf{y}) - f(\mathbf{x})) \cdot (Df((1-\xi)\mathbf{x} + \xi\mathbf{y}) \cdot (\mathbf{y} - \mathbf{x})).$$

Applying the Cauchy-Schwarz inequality to the expression on the right hand side one thus obtains

$$|f(\mathbf{y}) - f(\mathbf{x})|^2 \leqslant |f(\mathbf{y}) - f(\mathbf{x})| \, ||Df((1-\xi)\mathbf{x} + \xi\mathbf{y})|| \, |\mathbf{y} - \mathbf{x}|$$

and hence

$$|f(\mathbf{y}) - f(\mathbf{x})| < k \, |\mathbf{y} - \mathbf{x}| \,.$$

Corollary 1. Let $C \subset \mathbb{R}^n$ be non-empty, compact, convex and $f : C \to C$ continuously differentiable with $||Df(\mathbf{x})|| < 1 \; \forall \mathbf{x} \in C$. Then

$$\exists! \mathbf{x} \in C \text{ such that } f(\mathbf{x}) = \mathbf{x}.$$

Proof. Apply Lemma 1 and the Banach fixed point theorem.

The problem with Corollary 1 is to determine a rigorous upper bound for the operator norm of the derivative. Definition 1 defines what the value is, but provides no direct way of

computing it for a given matrix. To that end, recall the following standard result of matrix theory [5]:

$$||A|| = \sqrt{\lambda_{\max}(A^*A)}$$

where again $A \in \mathbb{C}^{n \times n}$, and A^* denotes the adjoint matrix, i. e. the transpose conjugate of A, and $\lambda_{\max}$ denotes the largest eigenvalue of a matrix.

Unfortunately, while A^*A can be computed easily, determining the largest eigenvalue of a matrix rigorously is non-trivial in general. This is particularly true in higher dimensions. For the special case of dimension 2, however, the eigenvalues can be computed by an analytic expression. So in this special case this approach is rather straightforward and provides an analytic way to compute the matrix norm and verify the contraction property.

Since direct calculation of of the operator norm is not straightforward, another option is to use a more readily computed quantity to provide an upper bound of $||A||$. Consider, for example, the following upper bound:

$$||A|| \leqslant \sqrt{n}||A||_\infty \tag{2}$$

where $||A||_\infty$ indicates the maximum absolute sum of the rows of A, i. e.

$$||A||_\infty = \max_{i=1,\dots,n} \sum_{j=1}^n |a_{i,j}|.$$

Proof. For any $\mathbf{x} \in \mathbb{R}^n$ we have

$$|A\mathbf{x}|^2 = \sum_{i=1}^n \left(\sum_{j=1}^n a_{ij}x_j \right)^2 \leqslant n \max_{i=1,\dots,n} \left(\sum_{j=1}^n |a_{ij}||\mathbf{x}| \right)^2 = n|\mathbf{x}|^2 \left(\max_{i=1,\dots,n} \sum_{j=1}^n |a_{ij}| \right)^2.$$

Clearly, $||A||_\infty$ is easily computed in a verified setting for any matrix A. Unfortunately, in general this inequality is a rather crude upper bound of the operator norm $||A||$.

Using inequality (2), it is possible to obtain a much more accurate condition to test for the contraction property of f under certain circumstances. Since we are only interested in showing the uniqueness of a fixed point, it is irrelevant in what coordinate system the map is a contraction. The following approach to compute the operator norm $||A||$ does so by changing into a more suitable coordinate system.

Let $T : C \to \mathbb{R}^n$ be an invertible, smooth coordinate transformation from C to $T(C)$. Then the following are equivalent:

- f has a unique fixed point in C;

- $F = T \circ f \circ T^{-1}$ has a unique fixed point in $T(C)$.

This is clear since for $x \in C$ such that $f(x) = x$ one finds that $T(x)$ satisfies $F \circ T(x) = T(x)$ and vice versa for $y \in T(C)$ such that $F(y) = y$, $T^{-1}(y)$ satisfies $f \circ T^{-1}(y) = T^{-1}(y)$.

Let the derivative of F be decomposed into a purely diagonal part L and a purely off-diagonal part N such that

$$DF = L + N.$$

Then applying the triangle inequality yields

$$||DF|| \leqslant ||L|| + ||N||.$$

The operator norm $||L||$ is simply the largest absolute value of the diagonal entries of L, while $||N||$ can be estimated by inequality (2).

Let H denote the convex hull of $T(C)$. By the assumptions of Corollary 1, Theorem 1 applies and yields at least one fixed point of f in C and hence a fixed point of F in $T(C) \subset H$. If $||DF|| < 1$ on H, Lemma 1 applied to the function F on H establishes the uniqueness of that fixed point and thus the uniqueness of the corresponding fixed point of f in $C \subset T^{-1}(H)$.

Note that if T can be chosen such that the derivative of F on $T(C)$ is approximately diagonalized, the entries in N are small compared to those in L, and the overestimation introduced by the estimation of $||N||$ is negligible compared to $||L||$ and the bound in $||DF||$ will be a rather tight one. Furthermore, if $T(C)$ itself is already convex, there is also no additional overestimation introduced by evaluating the derivative over the convex hull of $T(C)$.

Such a coordinate transformation T applicable in the neighborhood of some point $x \in C$ can be found relatively easily if $Df(x)$ is diagonalizable at x. Then T can be chosen to be an approximation of the matrix of eigenvectors. In that case, also $T(C)$ will be convex if C was convex, as T is a linear transformation.

To obtain a sharp bound, the derivative is only required to be approximately diagonal. It is thus not necessary to compute eigenvalues or eigenvectors rigorously. Only the matrix inversion T^{-1} has to be performed rigorously.

Preconditioning. While Theorem 1 is well suited for numerical computations, it poses one major problem. In the case of a hyperbolic, elliptic, or expanding fixed point, Theorem 1 does not apply directly at all, as there is generally no set C around the fixed point that is mapped into itself.

Even if applied directly to an attractive fixed point, the self mapping property may not succeed. One approach lies in performing all computations in a coordinate system aligned with the eigenvectors to overcome this problem.

There is, however, a much more powerful method to transform the map into a form that is applicable to Theorem 1 based on the well known Newton Method [5] used to approximate roots of differentiable functions.

Consider the operator

$$\mathcal{C}_A : f \mapsto A \cdot (f - \mathrm{id}) + \mathrm{id},$$

where A is a regular matrix and id represents the identity operator.

Since $f(x) - x = 0 \Leftrightarrow A\,(f(x) - x) = 0$ for regular A, we find that any fixed point of the map $\mathcal{C}_A f$ is also a fixed point of the map f and vice versa. We say that f and $\mathcal{C}_A f$ are equivalent with respect to their fixed points.

Assume x is a fixed point of f and hence $\mathcal{C}_A f$. We want to choose the matrix A in such a way that x becomes a strongly attracting fixed point of $\mathcal{C}_A f$. To that extend, consider the derivative of $\mathcal{C}_A f$:

$$D(\mathcal{C}_A f) = A \cdot (Df - I) + I.$$

If $Df(x) - I$ is a regular matrix, i. e. 1 is not a singular value of Df, then fixing $A = -(Df(x) - I)^{-1}$ yields

$$D(\mathcal{C}_A f)(x) = -I + I = 0.$$

That is, with this choice of A, $\mathcal{C}_A f$ has a strongly contracting fixed point at x. This is true independently of the type of fixed point the map f has at x. Furthermore, if $Df - I$ is not regular, it is possible to perturb it a little such that it becomes invertible. In that

case, the derivative of $\mathcal{C}_A f$ will not vanish completely, but if the perturbation was is small enough, $\mathcal{C}_A f$ will still be a strong contraction.

Thus, by applying Theorem 1 on $\mathcal{C}_A f$ and a non-empty, convex, compact set C around x, it is possible to obtain verification of the required self mapping property much more easily than by direct application to f.

The operator $\mathcal{C}_A$ is also very advantageous in the proof of uniqueness of a fixed point in C with Theorem 2. Since A is chosen such that the derivative of $\mathcal{C}_A f$ vanishes at the fixed point x, by continuity the ∞-norm of $D(\mathcal{C}_A f)$ is small on all of C, provided C itself is small enough. It is then possible to apply Lemma 1 to easily show uniqueness of the fixed point of $\mathcal{C}_A f$, and thus of f, in C by bounding $D(\mathcal{C}_A f)$ over C and utilizing inequality (2).

The effect of this preconditioning is shown in fig. 1. The dashed box has side length 10^{-4} in each direction and is centered around a period 15 point of a map of the Hénon family. The solid black object indicates the return of that box after 15 iterations. Without preconditioning, the resulting object is mostly linear with relatively small non-linear curvature. Clearly, the image does not map into the original box. In the preconditioned case, the same initial box is shrunk into a very highly non-linear object at least one order of magnitude less in size than the original box.
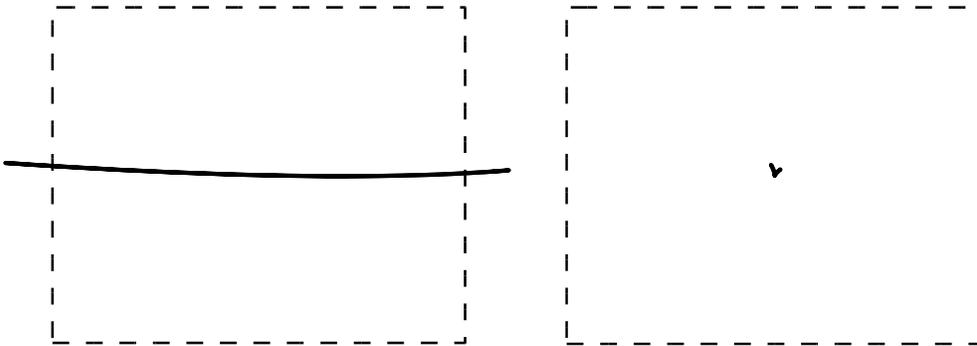
Figure 1: The effect of preconditioning, a box of size $10^{-4} \times 10^{-4}$ in a Hénon map and its 15th iterate without (left) and with (right) preconditioning

Implementation. The operator described above provides a versatile method to classify single boxes around a fixed point. In the following, we extend this to an algorithm for a global fixed point finder as described in the introduction. To easily satisfy the requirement of Theorem 1 that the sets must be convex, compact and non-empty, our implementation will only operate on sets that are interval boxes. Interval boxes are subsets of $\mathbb{R}^n$ given by a non-empty, finite interval in each component, i. e.

$$B = [\underline{b}_1, \overline{b}_1] \times [\underline{b}_2, \overline{b}_2] \times \cdot \times [\underline{b}_n, \overline{b}_n].$$

Clearly, interval boxes are convex and as long as $\underline{b}_i, \overline{b}_i \in \mathbb{R} \; \forall i = 1, \ldots, n$, they are compact, and if $\underline{b}_i < \overline{b}_i \; \forall i = 1, \ldots, n$ they are non-empty.

Let K be the interval box in which one wants to find the fixed points. The simplest method is to divide the search space into an even grid of fixed size boxes, and then checking each single small box for a fixed point using the classification presented in the previous section. However, even in two dimensional maps, that approach often is prohibitively expensive

in terms of computations required. In order to compute all fixed points in a box of size 1×1 to an accuracy of 10^{-4}, already $10^4 \cdot 10^4 = 10^8$ boxes must be checked.

Instead of fixing a grid, it is better to adaptively select the size of boxes to test in each region of the map based on the complexity of the map in that region. We use a stack-based "divide-and-conquer" style algorithm to successively remove parts of the search space that do not contain fixed points. Whenever no fixed point is present in the currently tested box, it is discarded. If there is no conclusive result for the current box either way, i.e. it may or may not contain a fixed point, the box is split into smaller boxes and analysis continues with one of the smaller boxes.

At this point, we intentionally leave it open how exactly the "size" of an interval box is measured, as such a measure may depend on the application. Commonly employed measures include the volume of the interval box, or the length of the longest side. Furthermore, it is not necessary at this point to specify an algorithm for the "splitting" of a box, as long as the employed algorithm results in a reduced measure of the split boxes. A simple splitting algorithm is to divide one of the longest sides of the box in half.

To guarantee that the search for fixed points terminates, there are certain pre-prescribed thresholds relating to the size of the boxes under consideration. The first is the desired maximum size of the verified fixed point enclosures, denoted by Ω_1. All boxes that contain or may contain fixed points will be at most of this size. The second constant, Ω_2, provides a lower limit on the size of boxes that are considered. Once the size of a box falls below this threshold, it will not be split any further, independently of its current status. We will discuss these limits in more detail below.

Algorithm. The basic algorithm for the verification procedure with automatic box size control is as follows:

1. Start with initial box K on the stack.

2. Take the top box S off of the stack and classify the box:

   - if S is verified to contain no fixed points, discard it;
   - else, if the size of the S is larger than Ω_1, split S and push the two resulting pieces onto the top of the stack;
   - else, if there is a unique fixed point in S, store S in the list of found unique fixed points;
   - else, if the size of S is larger than Ω_2, split along the longest side of S and push the two pieces onto the top of the stack;
   - otherwise, place S in the list of non-unique fixed points if existence was verified, else in the list of undecided boxes.

3. Continue with 2 until there are no more boxes on the stack.

In practice, the classification of a box S is performed in several steps, checking after each if further tests are necessary. The following gives the order in which verification is performed:

1. Compute $f(S)$ in Taylor Model arithmetic. If $f(S) \bigcap S = \emptyset$, stop as there is certainly no fixed point of f in S.

2. Compute an approximate $A = (Df(x) - I)^{-1}$ in floating point arithmetic, for some $x \in S$. If the inverse does not exist or is very badly conditioned, set A to a more suitable "approximate" inverse.

3. Compute $\mathcal{C}_A f(S)$ in Taylor Model arithmetic. If $\mathcal{C}_A f(S) \not\subset S$, stop as existence is not verified.

4. Compute $D\left(\mathcal{C}_A f\right)(S)$ in Taylor Model arithmetic using a provided analytic expression for Df and subsequently test the bound on the operator norm of the derivative to be less than one to verify uniqueness.

Analysis. The use of Taylor Models to carry out the verified computations makes the computation of the matrix A in the construction of the operator $\mathcal{C}_A$ particularly easy. By evaluating f over an interval box, very good approximations of the derivatives at the center point of the box can be obtained very easily from the first order coefficients of the resulting Taylor Model. Due to the differential algebraic nature of Taylor Models, no separate computation is necessary.

The tests for inclusion of a Taylor Model in the original set are done by computing interval bounds on each Taylor Model. In general, it is true that interval enclosures of Taylor Models are substantially larger than the Taylor Model itself. In this particular case, however, the use of interval bounds suffers from no additional overestimation beyond that of the quality of the generated interval enclosure. This is because the initial box S is itself an interval box. Thus for $\mathcal{C}_A f(S)$ to lie within S, each coordinate bound of $\mathcal{C}_A f(S)$ must lie in the corresponding interval of S.

The use of Taylor Models for these computations is essential for another reason. For the contraction operator $\mathcal{C}_A$ to work properly, it is required to have strong control over the dependency problem. This is because $\mathcal{C}_A$ is constructed above such that the first derivative is nearly cancelled out. This cancellation, however, can only occur if it is computed symbolically, e. g. using Taylor Models. With mere interval evaluation, the dependency problem caused by this cancellation would lead to vast blowup of the resulting interval, and thus render the operator inefficient. The reader familiar with verified global optimization will find the above algorithm very familiar. Most verified global optimizers, such as for example COSY GO [6,7], use very similar methods to successively examine the entire search space and prune regions where there is no extremum.

This global fixed point finder goes beyond the minimal task of global optimization of enclosing the optimum by not only producing a list of enclosures of fixed points, but also verifying and classifying them in the same step, thus typically guaranteeing the existence and in most cases also the uniqueness of a fixed point in each resulting box.

Constants. As described above, there are two constants controlling the behavior of the algorithm with respect to box sizes. The first constant, Ω_1, is the maximum size of boxes stored in the results list. Any boxes larger than this size are either discarded because they do not contain fixed points, or are split into smaller boxes to be analyzed. To the user, this parameter specifies the desired accuracy of the fixed point enclosures.

The second constant, Ω_2, is the minimum size below which a box is never split any further. Boxes that have been split blow this size, and could not be classified yet, will be added to the "unclassified" list. From the user's perspective, this value is a cutoff value when the algorithm should give up attempts to classify a box.

Therefore, resulting boxes in any of the three lists (unique fixed points, fixed points, and undecided) will always range in size between Ω_1 and Ω_2.

The ideal choice of these constants is problem dependent and can greatly affect the performance of the overall algorithm. In general, of course, $\Omega_1 \geqslant \Omega_2$ should hold. In practice, however, there is a lower limit for the size of Ω_1 below which the algorithm becomes significantly less effective due to limited floating point precision. In our experiments, a good choice for the constant Ω_1 in double precision computations is about 10^{-7}.

Dependency Problem and Order Loss. Since the objective functions that need to optimized in our problem of finding period n periodic points consists of iteratively evaluating the objective function n times, the resulting code for a single evaluation is quite extensive and very susceptible to the dependency problem. In this context, the use of Taylor models provides a significant advantage because of their ability to control the dependency problem [8].

However, even when using Taylor models, the limited accuracy of floating point arithmetic leads to a secondary but important effect, namely an effective loss of computation order when the coefficients of Taylor Models become too small. The coefficients of a converging Taylor Model scale with the order of their monomial. In particular, the error bound of a Taylor Model of order n scales as $O(x^{n+1})$ with the width of the Taylor Model [9].

This scaling law of Taylor Models holds true as long as the limited precision of floating point numbers is not considered. However, in an IEEE 754 [10] compliant double precision floating point environment, such as utilized by COSY INFINITY, a floating point number is only stored with about 16 significant decimal digits. This finite precision requires careful bookkeeping of the truncation errors during Taylor Model operations. Great care has been taken in implementing Taylor Models to bound all those errors in the remainder bound of the Taylor Model.

However, the round-off error of each floating point operation does not scale according to the above Taylor Model scaling law. The round-off error for a given coefficient of the Taylor Model resulting from a Taylor Model operation only scales as the coefficient itself. So for linear coefficients, the round-off error only scales linearly, while the round-off error for the constant part does not scale at all.

Under most circumstances, this poses no problem as the absolute value of the round-off error is very much smaller than the error due to the truncation of the Taylor series. However, as the coefficients of a Taylor Model become smaller, the size of the truncation error is greatly reduced, and the non-scaling round-off errors become a relatively larger contribution to the error.

Eventually, the higher order terms of the Taylor Model get smaller than the round-off errors. Effectively then those orders are lost, in the sense that keeping these terms in the Taylor Model does not increase the accuracy of the computation. We call this effect *order loss*.

Consider as an example an the function

$$f(x) = (1 + x + x^2)/3.$$

Evaluate this function with a Taylor Model box of width $2 \cdot \epsilon$ around 0 of the form $T = \epsilon \cdot x$. Since the operations are performed in double precision floating point arithmetic, the round-off error for the resulting constant part $(1/3)$ is on the order of 10^{-16} as $1/3$ is not a finite binary fraction. The round-off error of $\epsilon/3$ and $\epsilon^2/3$, respectively, is on the order of $10^{-16} \cdot \epsilon$ and $10^{-16} \cdot \epsilon^2$.

Let now $\epsilon = 2^{-10} \approx 10^{-3}$. That way ϵ and ϵ^2 are exact floating point numbers, and the only round-off error occurs when dividing by 3. Floating point round-off errors for each of the resulting coefficient are on the order of 10^{-16}, 10^{-20}, and 10^{-24} respectively and thus taken all together the error bound of the resulting Taylor Model is of order 10^{-16}, while the second order coefficient is $1/3 \cdot 2^{-20} \approx 3 \cdot 10^{-7}$.

Consider now the case where $\epsilon = 2^{-30} \approx 10^{-9}$. Again ϵ and ϵ^2 are exact floating point numbers, and the only round-off error occurs when dividing by 3. The round-off errors for each coefficient are now on the order of 10^{-16}, 10^{-25}, and 10^{-34} respectively. Clearly, since the round-off error of the constant part does not scale with ϵ, the resulting error bound is still on the order of 10^{-16}. However, this time the second order coefficient, $1/3 \cdot 2^{-60} \approx 3 \cdot 10^{-19}$ is much smaller than the error bound. Thus the contribution of the second order coefficient to the value of the Taylor Model over the domain $[-1, 1]$ is negligible relative to the error due to round-off.

Carrying out the above computation using only first order Taylor Models would yield the same quality result as the second order computation. Thus effectively the second order is lost, and the error scales no better than that of a first order Taylor Model. As ϵ gets even smaller, the Taylor Model asymptotically behaves no better than a classic interval.

The only way to prevent this order loss is to perform the computation using a higher precision for the coefficients. That way the round-off error of the constant part is reduced, and higher order computations remain viable for smaller values of ϵ. Of course for any finite precision, there exists an ϵ such that the Taylor Model behaves no better than an interval.

Note that this is a markedly high-order effect. The closest analog of order loss in interval arithmetic is practically never observed because it only occurs with very narrow intervals. Consider an interval in which both the lower and upper bound are stored as double precision numbers. Since double precision numbers are discreet, there is a minimum width for an interval. Once that minimum width has been reached, attempts to reduce the width of the interval will not lead to a smaller result.

In practice that means an interval around the number 1 cannot be smaller than about 10^{-16}, which for all practical purposes does not pose a problem. In Taylor Models, however, this effect is much more pronounced in the higher order terms, because of the Taylor Model scaling law.

Periodic Points in Accelerator Maps. We implemented the above algorithm for a fixed point finder in two dimensions using the rigorous COSY INFINITY [2] programming environment. We then applied it to an accelerator transfer map representing the Tevatron accelerator.

The transfer map is an 11th order map, computed by Pavel Snopok using the COSY INFINITY beam physics package [11]. The model of the Tevatron used is a slightly simplified model consisting of all dipole, quadrupole and sextupole bending magnets, but omitting certain corrective beam line elements, such as electrostatic separators and solenoids. This omission still produces a very close approximation of the real dynamical system. Moreover, it results in a symmetric map in which the x-a plane is invariant.

For this application, we restrict ourselves to the two dimensional motion in the invariant $x - a$ plane. If there is chaos in the full system, it is very likely to also be present in the two dimensional subspace. For this purpose, we extracted from the full four dimensional map the two dimensional map of the $x - a$ plane into itself.

Using the particle tracking facilities of COSY's beam physics package, we produced a tracking picture of the system using evenly spaced initial particle coordinates (see fig. 2, top).

0.600E-01

0.900E-01

0.600E-01

0.900E-01
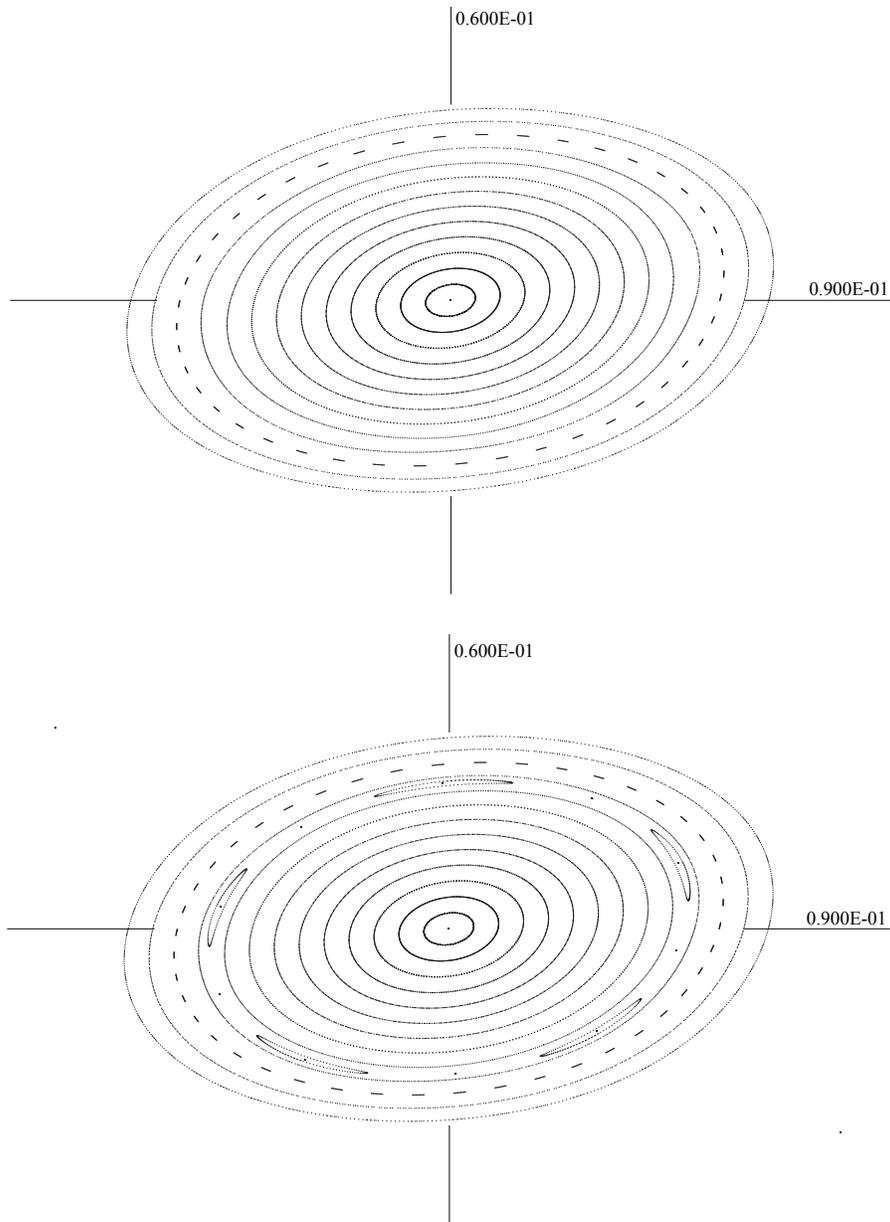
Figure 2: Tracking of evenly spaced particles through the transfer map in the $x - a$ plane (1400 iterations, 5 mm spacing between rings) (top) and location of the periodic points identified (bottom)

Between the rings are two sets of period 5 points, the elliptic ones with their associated islands around them. Two period 8 points are in the top left and bottom right corner.

Periodic points found for periods from 1 through 20
in the region [-0.09, 0.09]×[-0.06,0.06]

Period	Unique boxes	Area (Ω_1)	Time, <sec>	Periodic points	Tune
1	1	$4 \cdot 10^{-19}$	0	1	0.592
2	1	$4 \cdot 10^{-19}$	0	0	1
3	1	$4 \cdot 10^{-19}$	0	0	$2/3 \approx 0.666$
4	1	$4 \cdot 10^{-19}$	0	0	$3/4 = 0.75$
5	11	$4 \cdot 10^{-19}$	6	10	$3/5 = 0.6$
6	1	$4 \cdot 10^{-19}$	0	0	$4/6 \approx 0.666$
7	1	$4 \cdot 10^{-19}$	0	0	$5/7 \approx 0.714$
8	3	$4 \cdot 10^{-19}$	6	2	$5/8 = 0.625$
9	1	$4 \cdot 10^{-19}$	1	0	$6/9 \approx 0.666$
10	11	$4 \cdot 10^{-16}$	13	0	$6/10 = 0.6$
11	1	$4 \cdot 10^{-16}$	1	0	$7/11 \approx 0.636$
12	1	$4 \cdot 10^{-16}$	2	0	$8/12 \approx 0.666$
13	1	$4 \cdot 10^{-16}$	107	0	$8/13 \approx 0.615$
14	1	$4 \cdot 10^{-16}$	2	0	$9/14 \approx 0.643$
15	11	$4 \cdot 10^{-16}$	21	0	$9/15 = 0.6$
16	3	$4 \cdot 10^{-12}$	18	0	$10/16 = 0.625$
17	1	$4 \cdot 10^{-12}$	4	0	$11/17 \approx 0.647$
18	1	$4 \cdot 10^{-12}$	359	0	$11/18 \approx 0.611$
19	1	$4 \cdot 10^{-12}$	10	0	$12/19 \approx 0.632$
20	11	$4 \cdot 10^{-12}$	33	0	$12/20 = 0.6$

Estimating the region of stability to be within the region shown, $[-0.09, 0.09] \times [-0.06, 0.06]$, we started the periodic point finder to work on that region.

Table shows the number of enclosures containing a unique periodic point found for periods up to 20. As all COSY beam maps, this map is origin preserving since it represents the reference trajectory. Hence for every period, there is at least one periodic point identified, namely the origin. Similarly, periodic points of period 5 are also listed for periods 10, 15, and 20. But since all of the periodic point enclosures contain exactly one periodic point, and we know that periodic points of lower period must be present, we can subtract out those points of lower period from the results. The count of periodic points without those duplicate points is given in the second to last column.

This computation rigorously proves that there are only periodic points of periods 5 and 8 within the region of interest. Note that for the period 5 points, there are two separate sets of 5 points each, while for the period 8 points only two were found. The missing 6 points of their orbit lie outside the search area and thus are not identified. Fig. 2 (bottom) illustrates the location of all periodic points identified up to period 20.

Due to the splitting algorithm we chose in our implementation, always cutting along the longest side of a box, the resulting periodic point enclosures are roughly square. Their area is less than the parameter Ω_1 we specified, given in the area column. As is expected due to the broken scaling law it was necessary to increase the targeted area as the period increases to guarantee successful verification. Note, however, that even the enclosures at period 20 still are accurate up to an error of about 10^{-6}.

The time given is the CPU time consumed for the search rounded to the nearest full sec-
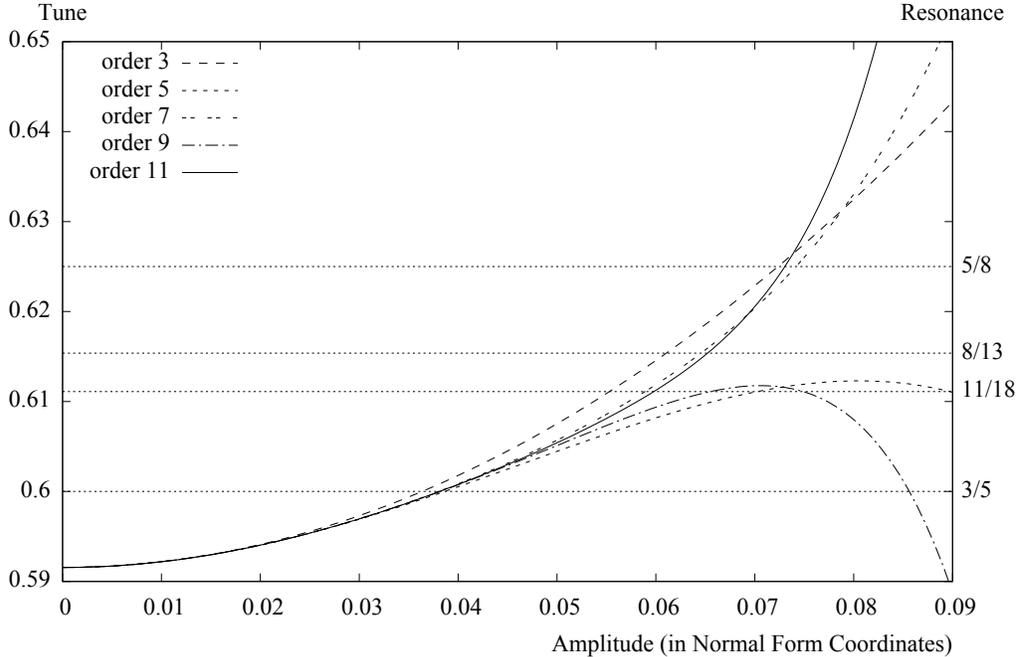
Figure 3: Polynomial expansions of tune shifts computed to orders 3, 5, 7, 9, and 11 along with the first low order resonances

ond and measured running on a single core of our Intel Xeon X5677 system at 3.5 Ghz. It is interesting to note the vast differences in execution times between different periods. Particularly periods 13 and 18 take significantly longer than all other periods. This phenomenon will be analyzed in more detail in the next section.

Tune Shifts. The Tevatron map analyzed by us exhibits a linear tune of about 0.5915. The accelerator was intentionally designed such that the linear tune is close to the fundamental resonance $3/5 = 0.6$. Normal form theory predicts a shift of the tune as the distance form the reference orbit increases. Those tune shifts can help explain why we found period 5 periodic points.

Normal form transformations based on differential algebra methods [12] applied to our beam transfer map $\mathcal{M} : \mathbb{R}^2 \to \mathbb{R}^2$ provide a coordinate transformation $\mathcal{T} : \mathbb{R}^2 \to \mathbb{R}^2$ such that the map $\mathcal{T} \circ \mathcal{M} \circ \mathcal{T}^{-1}$ takes elliptic orbits into circular orbits around the origin up to a given order, and the tune shift $t(r)$ only depends on the radius of the circle [13].

When the tune shift $t(r)$ attains a rational value of the form n/k with $n, k \in \mathbb{N}$, periodic points of period k are expected to appear in the map in normal form coordinates. Consequently, the original map $\mathcal{M}$ also should exhibit periodic points of the same period.

We used the COSY INFINITY beam physics code to compute an expansion of the tune shift in normal form coordinates [12] to various orders. In fig. 3, those tune shifts are plotted as a function of the radius in normal form coordinates. Also marked are several interesting resonances.

The wild oscillations of the function at different orders for large radii indicate that the polynomial expansion of the tune shift does not converge very well for values above ap-

proximately 0.06. However, for values below about 0.05 the plots are in good agreement, suggesting good convergence in that region.

The last column of table shows the nearest fraction $n/p, n \in \mathbb{N}$ for each period p that is larger than the tune of the Tevatron, 0.5915. We chose to list larger values as the tune appears to be shifting upwards in the tune shift plots. As mentioned before, the resonance for the period 5 points is by design very close to the tune of the Tevatron. This is because then automatically other low order resonances occur significantly further from away the Tevatron tune.

The next two closest resonances, occurring for periods 13 and 18, are $8/13 \approx 0.615$ and $11/18 \approx 0.611$. Interestingly enough, those two periods were also the two periods taking the longest for the global periodic point finder to treat, taking up to 12 times longer than all other periods. This indicates that those periods are "closer" to produce periodic points, i. e. points do return to within very close proximity of their origin. However, our periodic point finder has shown that none actually return exactly to themselves (except the origin). This complicates the search for our algorithm and results in significantly longer execution times.

From the normal form tune shift plots in fig. 3 we can determine the radius for which the tune shift reaches $3/5 = 0.6$. In all orders it is consistently located at about 0.0387. At this tune in the normal form we would expect period 5 points to appear. To compare this prediction with the results of the periodic point finder, we transformed the coordinates (x, a) of the center of one of the identified period 5 enclosures into normal form coordinates $(X, A) = \mathcal{T}(x, a)$ and computed its radius $R = \sqrt{X^2 + A^2}$.

The point we chose for this is $x \approx 0.00137$ and $a \approx -0.0295$, which, in COSY's normal form coordinates, becomes $X \approx 0.00631$ and $A \approx -0.0375$, and thus has a radius $R \approx 0.0381$. As can be seen, both the radius predicted from the normal form tune shift plot as well as the periodic point identified by the periodic point finder are in agreement within about 2% of each other.

Unfortunately, this is the only valid prediction that can be made from the tune shift plot. In all of the plots, the next tune, 11/18, seems to be crossed. However, the periodic point finder proved that there are no periodic points of period 18 related to this tune within the search region.

The reason for this is the slow convergence of the normal form transformation due to the closeness of the linear tune to the 3/5 resonance. The tune shift expansions shown above do not converge for the values at which 11/18 would be obtained.

Converting one of the period 8 points we identified by the same normal form coordinate transformation as used for the period 5 point above, we find that its amplitude would be about 0.238. This is clearly far beyond the convergence of the normal form expansion, hence its existence could not be predicted by the tune shift plots.

This comparison highlights a problem associated with the normal form methods applied to this type of problem. Since accelerators are intentionally designed to operate near one of the fundamental resonances, the convergence of the tune shift expansion is slow and the region in which reliable predictions can be made is rather small. Unfortunately, with non-verified computations it is not clear a priori how large that region is. Verified tools, such as this periodic point finder, can help in identifying rigorously such regions.

Manifolds. For further analysis, we concentrate on the hyperbolic period 5 points identified in the stable region of the map. In order to show the existence of chaos, we numerically computed the one dimensional stable and unstable manifolds of one of these points. If there are homoclinic intersections, i. e. intersection between the stable and
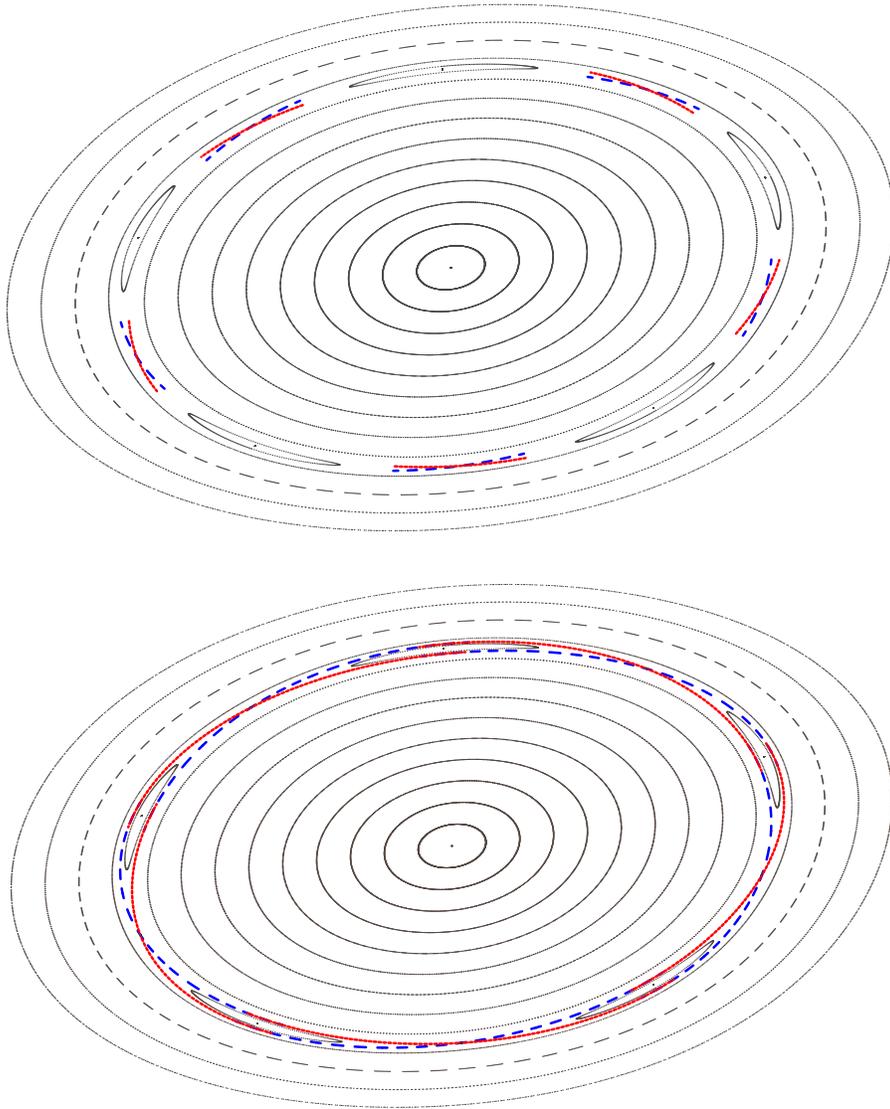
Figure 4: Top: the initial polynomial approximation of the solid stable (red) and dashed unstable (blue) manifolds of the hyperbolic period 5 points in the tracking picture; bottom: the full manifold braid after 50 turns

unstable manifolds, this gives rise to chaotic motion.

Using methods developed by Johannes Grote e. a. [14-16] we compute a 20th order polynomial approximation of the manifolds at the period 5 point in the interval box $(x, a) = ([0.00137155595, 0.00137155598], [-0.0295197926, -0.0295197924])$. The verified enclosures of the local manifolds generated by this method were small pieces of length on the order of 0.003 and a thickness of about 0.0001.

There are two reasons for those relatively poor verified enclosures. For one, the initial enclosure of the periodic point is not sharp enough for a good estimate. While an enclosure of approximately 10^{-9} in width may be very good for practical applications, it is too big an enclosure for the verified manifold generator. It causes any rigorous function evaluation at the periodic point to automatically pick up an error bound of that size, independently of any other errors added by the computation. Due to this, even the non-verified polynomial manifold approximation is only mapped into itself with an error on the order of about 10^{-9}.

Secondly, evaluating the full polynomial map and its derivative in Taylor Model arithmetic over such a relatively large box, introduces considerable additional overestimation, especially for the derivative. This forces the verification algorithm to increase the error bound on the manifold enclosure, and together these two effects cause the relatively poor result.

To obtain much sharper enclosures of the manifolds that can be used for further, rigorous analysis, a higher precision enclosure of the periodic point is necessary. We believe that such a change will drastically increase the accuracy of the resulting manifold enclosures.

Meanwhile, we will carry out our analysis using the existing, non-verified DA approximations of the manifolds. The initial manifolds are shown in fig. 4, top, for a parameter domain of $[-1, 1]$. The manifolds for the other points in the orbit are obtained by simply iterating the manifolds for one of the points forward 4 more times.

In order to extend the length of the manifold, we iterate the initially generated manifold pieces through the map for the unstable manifold, or the inverse of the map for the stable manifold. The inverse of the map is here computed using non-verified DA inversion techniques.

Computation of the eigenvalues of five iterations of the Tevatron transfer map at the periodic points yield eigenvalues $\lambda_s = 0.969$ and $\lambda_u = 1.032$. Since the eigenvalues give the growth rate of the local manifold around the periodic points, this means that a large number of iterations is necessary to extend the length of the manifold.

Furthermore, with higher iterations convergence of the single polynomial approximation will become worse. To overcome this problem, we cover the manifold by several polynomials, each on the domain $[-1, 1]$, which are expanded locally, instead of around the periodic point. If after an iteration through the map the highest order coefficients of one of those polynomials become larger than 10^{-12}, we split this polynomial in the middle of the domain into two pieces, and reparametrize each back onto the domain $[-1, 1]$. This yields a list of polynomials that each cover a piece of the manifold locally. This algorithm is a non-verified version of the verified manifold iteration algorithm described in [14].

Fig. 4, bottom, shows the resulting manifold braid after 50 iterations. Each of the manifolds is covered by 6 polynomials. The plot suggests that the manifolds between the periodic points might form heteroclinic connections, i. e. the stable manifold of one point is also the unstable manifold of its neighbors, forming the braid-like pattern seen in the picture.

With very sharp, rigorous enclosures of the manifolds, it may be possible to show that the manifolds of one periodic point really do intersect the manifolds of its neighbor. If that

were the case, the connections certainly would not be heteroclinic. At this point, however, we are unable to decide the question whether these manifolds form homoclinic intersections or heteroclinic connections conclusively.

Conclusion. A fast, rigorous, global periodic point finder has been implemented. The use of Taylor Models has been shown to produce fast results even for high period points in numerically difficult maps.

The main restriction to obtaining sharper enclosures and higher periods is the limited precision of the underlying double precision floating point numbers. With the advent of high precision Taylor Models, which are currently in development, the order loss associated with limited precision can be contained and even better results will be obtained.

Applying these techniques to dynamical systems such as beam transfer maps allows the automated analysis of those maps for regions of stability and aides in identifying regions of potential chaos. The rigorous character of the method presented augments the qualitative results obtained from existing methods of analysis.

While currently we cannot decide whether homoclinic intersections occur in this map, we believe that here, too, high precision Taylor Models will allow us a more detailed, rigorous analysis. Further analysis of the mathematical consequences of the symplecticity of the transfer map may reveal an analytical explanation for the existence of heteroclinic connections between periodic points, although the authors are not aware of any such results in the literature.

References

1. Berz M. From Taylor series to Taylor models. *AIP CP*, 1997, vol. 405, pp. 1–20.

2. Makino K., Berz M. COSY INFINITY version 9. *Nuclear Instruments and Methods*, 2006, vol. 558, pp. 346–350.

3. Schauder J. Der Fixpunktsatz in Funktionsräumen. *Studia Mathematica*, 1930, vol. 2, pp. 171–180.

4. Banach S. Sur les opérations dans les ensembles abstraits et leurs application aux équations intégrales. *Fund. Math*, 1922, vol. 3, pp. 7–33.

5. Bronstein I. N., Semendjajew K. A., Musiol G., Mühlig H. Taschenbuch der Mathematik. 5th edition. Verlag Harri Deutsch, 2001, pp. 269, 652.

6. Makino K., Berz M. Range bounding for global optimization with Taylor models. *Transactions on Computers*, 2005, vol. 4, no. 11, pp. 1611–1618.

7. Makino K., Berz M. Rigorous global optimization of system parameters. *Vestnik of S-Petersb. University*, ser. 10: Applied mathematics, computer science, control processes, 2014, issue 2, pp. 61–71.

8. Makino K., Berz M. Efficient control of the dependency problem based on Taylor model methods. *Reliable Computing*, 1999, vol. 5, no. 1, pp. 3–12.

9. Berz M., Hoffstätter G. Computation and application of Taylor polynomials with interval remainder bounds. *Reliable Computing*, 1998, vol. 4, no. 1, pp. 83–97.

10. IEEE standard for binary floating-point arithmetic: technical report IEEE Std 754-1985. The Institute of Electrical and Electronics Engineers, 1987. Reprinted in New York: ACM SIGPLAN Notices, 1987, vol. 22, no. 2, pp. 9–25.

11. Snopok P. Optimization of Accelerator Parameters Using Normal Form Methods on High-Order Transfer Maps: PhD thesis. Michigan, USA: Michigan State University, East Lansing, 2007, URL: http://bt.pa.msu.edu/snopokphd.

12. Berz M. Differential algebraic formulation of normal form theory. *Proc. Nonlinear Effects in Accelerators*. Ed.: M. Berz, S. Martin, K. Ziegler. London: IOP Publishing, 1992, pp. 77–86.

13. Berz M. High-order computation and normal form analysis of repetitive systems. *Physics of Particle Accelerators.* Ed.: M. Month., New York: American Institute of Physics, 1991, vol. 249, pp. 456–489.

14. Wittig A., Berz M., Grote J., Makino K., Newhouse S. Rigorous and accurate enclosure of invariant manifolds on surfaces. *Regular and Chaotic Dynamics*, 2010, vol. 15, pp. 107–126. (also doi:10.1134/S1560354710020024).

15. Grote J., Berz M., Makino K., Newhouse S. Taylor model-based enclosure of invariant manifolds for planar diffeomorphisms and applications. *Proc. Applied Mathematics and Mechanics*, 2008, vol. 7, pp. 1022907–1022908.

16. Grote J. High-Order Computer-Assisted Estimates of Topological Entropy: PhD thesis. Michigan, USA: Michigan State University, East Lansing, 2008, URL: http://bt.pa.msu.edu/grotephd

Alexander Wittig – doctor of philosophy; e-mail: wittigal@msu.edu
Martin Berz – doctor of philosophy, professor; e-mail: berz@msu.edu