On the computation of invariant manifolds of fixed points

Homburg, A.J.; Osinga, H.M.; Vegter, Geert

Published in:
Default journal

Publication date:
1995

Z angew Math Phys *46* (1995)
(ZAMP)

0044-2275/95/020171-17 $ 1.50 + 0.20

On the computation of invariant manifolds of fixed points

By A. J. Homburg*, H. M. Osinga** and G. Vegter, Dept of Mathematics and Computing Science, P.O. Box 800, 9700 AV Groningen, The Netherlands (e-mail: gert@cs.rug.nl)

1. Introduction

Invariant manifolds play an important role in the qualitative analysis of dynamical systems. For classes of simple systems basins of attraction of periodic orbits are separated by stable or unstable manifolds. If a system depends on parameters it may exhibit non-local bifurcations due to a change in the geometric configuration of its invariant manifolds. Furthermore, certain forms of chaotic dynamics occur if a dynamical system has a homoclinic tangle, i.e. an intersection between the stable and unstable manifold of a periodic orbit. Detecting this kind of chaos therefore requires a rather precise knowledge of the invariant manifolds.

One of the first proofs of the *invariant manifold theorem* for hyperbolic periodic orbits (see Section 2 for terminology) was given by Perron, see [12], using a method related to *variation of constants*. An even earlier method by Hadamard, based on the *graph transform*, is of a more geometric flavor, see [4]. These methods and their generalizations, see e.g. [5, 13, 14, 15], are *constructive*, although usually far from being algorithmic. In these approaches an invariant manifold is obtained as a fixed point of an operator defined on a suitably defined metric space, whose elements represent (local) subspaces of $\mathbb{R}^d$. Hyperbolicity of the orbit guarantees that this operator is a contraction. A numerical method based on the graph transform is described in [7]. The variation of constants method is implemented in [16]. In [8] a shooting method is used to compute invariant manifolds of codimension one.

In this paper we derive a numerical algorithm for the computation of invariant manifolds of hyperbolic periodic orbits of diffeomorphisms using *invariant foliations*, a geometric tool for the qualitative study of dynamical

* Supported by NWO grant 611-307-018.
** Supported by NWO grant 611-306-523.

systems, see e.g. [5, 9, 10]. This approach has some features in common with both the graph transform and the method of variation of constants. A simple version of the algorithm, based on linear foliations, is presented in Section 2. In Section 3 we present a geometric derivation of the algorithm. We also show how the class of foliations can be slightly enlarged in order to make the algorithm very efficient. In Section 4 we discuss numerical issues, and we also give some examples. In Section 5 we generalize the algorithm, so that it also computes strong stable and unstable manifolds of pseudohyperbolic periodic orbits. Section 6 contains the proof of the correctness of our method.

2. Preliminaries and main result

First we introduce some terminology and basic properties of invariant manifolds and invariant foliations. For an in-depth treatment and more background material we refer to [9].

We mainly study discrete dynamical systems, viz diffeomorphisms, on a finite dimensional manifold in the neighborhood of a fixed point. Our method can easily be extended to continuous systems, see Section 4.

Since our study is local we consider a diffeomorphism $f\colon \mathbb{R}^d \to \mathbb{R}^d$ with fixed point $0 \in \mathbb{R}^d$. We say that $0 \in \mathbb{R}^d$ is a hyperbolic fixed point if no eigenvalue of $L := Df(0)$ has modulus 1. In this case there is a linear L-invariant decomposition $\mathbb{R}^d = E^s \oplus E^u$, such that the eigenvalues of $L^s := L \,|\, E^s$ are inside the unit circle in the complex plane and the eigenvalues of $L^u := L \,|\, E^u$ are outside the unit circle. The *stable* and *unstable manifold* of the map f at 0 are the sets W^s and W^u of points tending to 0 under forward and backward iteration, respectively. In other words, $x \in W^s$ iff $\lim_{n\to\infty} f^n(x) = 0$, and $x \in W^u$ iff $\lim_{n\to\infty} f^{-n}(x) = 0$. We use the equivalent definition of W^s and W^u as the set of points near 0 whose forward and backward orbits lie near 0. Note that E^s and E^u are the stable and unstable manifolds of L. According to the *invariant manifold theorem* for hyperbolic fixed points W^s and W^u are immersed submanifolds of $\mathbb{R}^d$ whose tangent spaces at 0 are E^s and E^u, respectively. If the map is C^k, $1 \le k \le \infty$, these manifolds are C^k as well.

For $p \in \mathbb{R}^d$ let $\mathscr{E}^s_p$ and $\mathscr{E}^u_p$ be the linear manifolds through p, parallel to E^s and E^u, respectively. A first, very simple, version of our algorithm takes as input a point p near 0, and computes the forward orbit of the point $\varphi(p) := \mathscr{E}^u_p \cap W^s$. The latter orbit is computed as the fixed point of a contractive operator, defined on the space Σ_p of *bounded sequences* of points near $0 \in \mathbb{R}^d$, whose initial point lies on $\mathscr{E}^u_p$. More precisely, this operator T maps a sequence $\{x_n\}_{n\ge 0} \in \Sigma_p$ onto the sequence $\{x'_n\}_{n\ge 0} \in \Sigma_p$, defined by:

$$x'_n = \begin{cases} [p, f^{-1}(x_{n+1})], & \text{if } n = 0, \\ [f(x'_{n-1}), f^{-1}(x_{n+1})], & \text{if } n > 0, \end{cases}$$

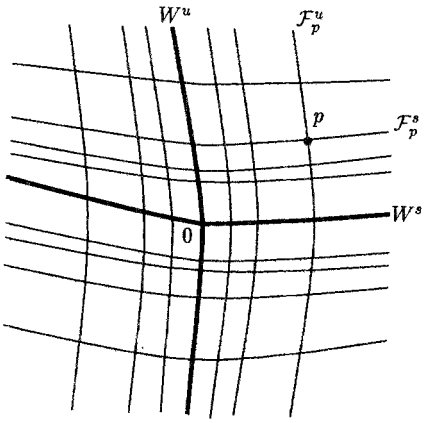

Figure 1
A pair of stable and unstable foliations.

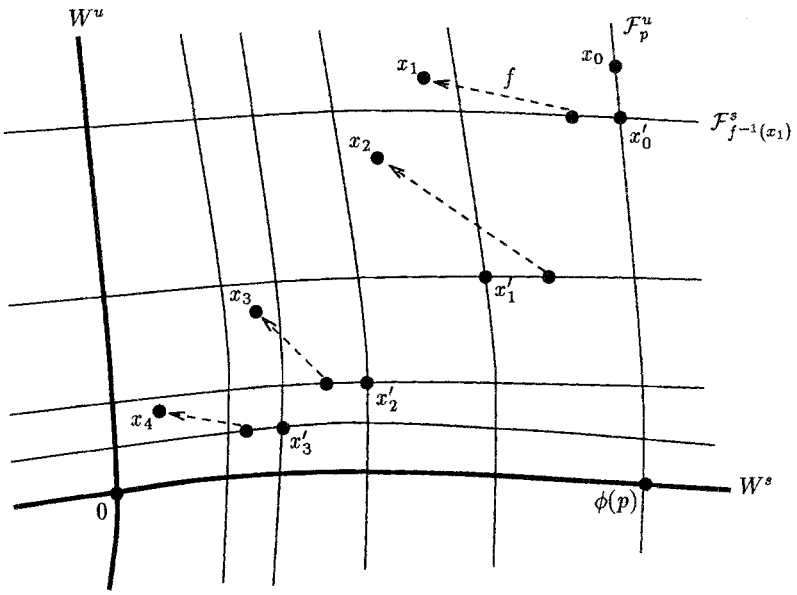

Figure 2
The contraction T.

where $[y, z]$ denotes the point $\mathscr{E}^u_y \cap \mathscr{E}^s_z$, for y, z near $0 \in \mathbb{R}^d$. The initial sequence to which T is applied may be any bounded sequence $\{x_n\}_{n \geq 0}$, such that $\mathscr{E}^u_p = \mathscr{E}^u_{x_0}$. An obvious choice is therefore the sequence defined by $x_0 = \mathscr{E}^u_p \cap E^s$, and $x_n = L(x_{n-1})$ $(= L^n(x_0))$, if $n \geq 1$.

The geometric background is illustrated in Fig. 2, for a more general situation. (If we replace the symbols $\mathscr{F}$ with $\mathscr{E}$ the figure also applies to the case we just described.)

In case f is linear, i.e. $f = L$, the iterates produced by the algorithm can be described explicitly. In this trivial case one can easily verify that the successive iterates converge to the orbit of the point $\varphi(p)$ on E^s $(= W^s)$.

It is easy to check that the orbit of $\varphi(p)$ is a fixed point of the operator T. To show that this fixed point is unique, we endow Σ_p with the sup-norm, and prove that T is a contraction. In other words:

Main result (special case). The operator $T: \Sigma_p \to \Sigma_p$ is a contraction. Its fixed point is the orbit of $\varphi(p) = \mathscr{E}_p^u \cap W^s$.

The contractive factor of T can be determined explicitly, see Theorem 3, the general version of the main result. It turns out that convergence of the algorithm does not require $\mathscr{E}_p^s$ and $\mathscr{E}_p^u$ to be *exactly* parallel to the spaces E^s and E^u. Exploiting this fact leads to a slightly different contraction, and hence to a slightly different version of the algorithm, that is *computationally more flexible*. In the next section we make these ideas more precise.

3. Geometric derivation of an algorithm

We state the algorithm in terms of a pair of invariant foliations $(\mathscr{F}^s, \mathscr{F}^u)$, defined on a neighborhood of the fixed point. The invariant manifolds W^s and W^u are special leaves of $\mathscr{F}^s$ and $\mathscr{F}^u$, respectively. Invariance of $\mathscr{F}^s$ boils down to $f(\mathscr{F}_p^s) = \mathscr{F}_{f(p)}^s$, locally near 0. Here $\mathscr{F}_p^s$ denotes the leaf of $\mathscr{F}^s$ through p, see Fig. 1. These foliations are not unique, but any pair $(\mathscr{F}^s, \mathscr{F}^u)$ of f-invariant foliations is near the pair of L-invariant linear foliations $(\mathscr{E}^s, \mathscr{E}^u)$. Furthermore each *leaf* of $\mathscr{F}^s$ and $\mathscr{F}^u$ is as smooth as f, but in general the *foliations* themselves are only continuous. We refer to [10], appendix 1, for an overview of these, and related, properties of invariant foliations.

For a point p near 0 the algorithm computes the point $\varphi(p)$, the point of intersection of $\mathscr{F}_p^u$ and $W^s = \mathscr{F}_0^s$. If we run the algorithm for sufficiently many points near 0, e.g. some large subset of E^s, we obtain an approximation of the stable manifold. A similar method can be used to compute an approximation of the unstable invariant manifold.

A first version of the algorithm is obtained by replacing the foliations $\mathscr{E}^s$ and $\mathscr{E}^u$ by $\mathscr{F}^s$ and $\mathscr{F}^u$ in the definition of the operator T, introduced in Section 2. However, there is a catch here: invariant foliations are hard to use computationally; leaves near the invariant manifolds are as complex as the invariant manifolds themselves. Even the representation of a single leaf in general will require complex data structures, leading to an unacceptable amount of storage and long computing times. Therefore this version of the algorithm, although simple in theory, is totally impractical. Fortunately, the

constraint on the foliations to be f-invariant can be relaxed. It is sufficient to use a pair of foliations that is near an f-invariant pair of foliations, like the pair of L-invariant foliations $(\mathscr{E}^s, \mathscr{E}^u)$, or any other pair that is computationally very convenient. This idea leads to the final version of the algorithm, presented in Section 3.2.

3.1. A first version of the algorithm

As we have seen in Section 2, we consider the space Σ_p of sequences of points near 0, starting on $\mathscr{F}_p^u$. To turn this space into a complete metric space we introduce a norm $\|\cdot\|$ on $\mathbb{R}^d$, such that $\lambda_s := \|L^s\| < 1$ and $\lambda_u^{-1} := \|(L^u)^{-1}\| < 1$. The metric d on $\mathbb{R}^d$ is defined by

$$d(x, y) := \max(\|\pi_s(x) - \pi_s(y)\|, \|\pi_u(x) - \pi_u(y)\|),$$

where π_s and π_u are the canonical projections $\mathbb{R}^d \to E^s$ and $\mathbb{R}^d \to E^u$, respectively. Note that both the norm $\|\cdot\|$ and the metric d depend on the operator L, although our notation does not make this explicit. Σ_p is a complete metric space with respect to the metric σ, defined by $\sigma(\boldsymbol{x}, \boldsymbol{y}) = \sup_{n \geq 0} d(x_n, y_n)$, for $\boldsymbol{x} = \{x_n\}_{n \geq 0}$, $\boldsymbol{y} = \{y_n\}_{n \geq 0} \in \Sigma_p$.

The definition of the contractive operator T arises from simple geometric observations. Recall that we want to compute an orbit in the stable manifold, which is a special element of the space Σ_p. Therefore we impose the following conditions on the T-image $\boldsymbol{x}' = \{x_n'\}_{n \geq 0}$ of a sequence $\boldsymbol{x} = \{x_n\}_{n \geq 0} \in \Sigma_p$:

1. Leaves on $\mathscr{F}^s$, through points of $\boldsymbol{x}'$, are closer to W^s than leaves of $\mathscr{F}^s$, through points of $\boldsymbol{x}$.
2. The family of leaves $\mathscr{F}_{x_n'}^u$, $n \geq 0$, is f-invariant.

Let $\varepsilon > 0$ be such that $\max(\lambda_s, \lambda_u^{-1}) + \varepsilon < 1$. Defining the distance between two leaves of $\mathscr{F}^u$ ($\mathscr{F}^s$) to be the distance between their points of intersection with W^s (W^u), we see that the canonical action of f on the space of leaves of $\mathscr{F}^u$ is a contraction, with contractive factor $\lambda_s + \varepsilon$, provided we consider leaves on a neighborhood of 0 (whose size depends on ε). Similarly the action of f^{-1} on the space of leaves of $\mathscr{F}^s$ is a contraction, with contractive factor $\lambda_u^{-1} + \varepsilon$. Therefore the leaf $\mathscr{F}_{f^{-1}(x_{n+1})}^s$ is closer to $W^s = \mathscr{F}_0^s$ than $\mathscr{F}_{x_{n+1}}^s$. Hence we take x_n' on the leaf $\mathscr{F}_{f^{-1}(x_{n+1})}^s$, thus satisfying condition 1. Condition 2 amounts to $x_n' \in \mathscr{F}_{f(x_{n-1}')}^u$, for $n > 0$. This discussion leads to the definition $T(\boldsymbol{x}) := \{x_n'\}_{n \geq 0}$, where

$$x_n' = \begin{cases} [p, f^{-1}(x_1)], & \text{if } n = 0, \\ [f(x_{n-1}'), f^{-1}(x_{n+1})], & \text{if } n > 0, \end{cases} \tag{T_1}$$

where $[y, z]$ denotes the point $\mathscr{F}_p^u \cap \mathscr{F}_z^s$, for y, z near $0 \in \mathbb{R}^d$. Also see Fig. 2. Note that we suppress the dependence of T on p and $(\mathscr{F}^s, \mathscr{F}^u)$ in our

notation. This way we achieve that the family of leaves $\mathscr{F}^u_{x_n}$, $n \geq 0$, is f-invariant, in the sense that $f(\mathscr{F}^u_{x_n}) = \mathscr{F}^u_{x_{n+1}}$. In particular $\mathscr{F}^u_{x_n} = f^n(\mathscr{F}^u_{x_0}) = \mathscr{F}^u_{f^n(x_0)}$, so, after the first iteration, we have $\mathscr{F}^u_{x_n} = \mathscr{F}^u_{x_n}$. It is easy to see that the forward orbit of the point $\varphi(p) = \mathscr{F}^u_p \cap W^s$ is a fixed point of T. Using the fact that the f-action on the space of leaves of $\mathscr{F}^u$ and the f^{-1}-action on the space of leaves of $\mathscr{F}^s$ are contractions, we see that T is a *contraction*, with contractive factor $\kappa := \max(\lambda_s, \lambda_u^{-1}) + \varepsilon$. Hence the forward orbit of $\varphi(p)$ is the unique fixed point of T. After n iterations the distance between the current sequence and this fixed point is at most a κ^n-fraction of the distance of the input sequence and the fixed point.

3.2. *An efficient algorithm*

We now relax the constraint that the foliations used in the definition of T are f-invariant. In this way we obtain a contraction that is still defined by (T_1), but whose implementation is simpler since we can choose foliations that are *computationally more convenient* than invariant foliations. This larger flexibility is paid for by the technical constraint that our foliations are of class C^1, cf. Section 6.

Definition 1. For $\gamma > 0$ the unstable cone through a point x near 0 is the subset $C^u_x(\gamma)$ of a neighborhood of 0 defined by

$$y \in C^u_x(\gamma) \text{ iff } \|\pi_s(x) - \pi_s(y)\| \leq \gamma \|\pi_u(x) - \pi_u(y)\|.$$

Similarly the stable cone through x is defined by

$$y \in C^s_x(\gamma) \text{ iff } \|\pi_u(x) - \pi_u(y)\| \leq \gamma \|\pi_s(x) - \pi_s(y)\|.$$

The families of cones $\{C^s_x(\gamma)\}$ and $\{C^u_x(\gamma)\}$, where x ranges over a neighborhood of 0, are called a *stable* and *unstable cone-field*, respectively. For $\gamma < 1$ the unstable cone-field is f-invariant in the sense that $f(C^u_x(\gamma)) \subset C^u_{f(x)}(\gamma)$, provided we restrict to some small neighborhood of 0, cf. [2], and Lemma 9. Similarly the stable cone-field is f^{-1}-invariant.

Definition 2. Let $\gamma > 0$ and let U be a neighborhood of the fixed point 0 in $\mathbb{R}^d$. A pair of C^1-foliations $(\mathscr{F}^s, \mathscr{F}^u)$ is called γ-skew on U, if for $x \in U$ we have $\mathscr{F}^s_x \cap U \subset C^s_x(\gamma)$ and $\mathscr{F}^u_x \cap U \subset C^u_x(\gamma)$.

A computationally convenient pair of γ-skew foliations is $(\mathscr{E}^s, \mathscr{E}^u)$, whose representation in a computation does not require complex data structures, also see Secton 4. In the first version of the algorithm, viz the successive application of the contraction T, we can replace the pair of f-invariant foliations with a pair of γ-skew foliations. To state the main

result we consider a neighborhood U of 0, and a pair $(\mathscr{F}^s, \mathscr{F}^u)$ of γ-skew foliations. Let

$$\kappa(\gamma, U) := (1-\gamma)^{-1}((1+\gamma)(\max(\lambda_s, \lambda_u^{-1}) + \varepsilon) + 2\gamma C) < 1, \tag{1}$$

where $\varepsilon := \mathrm{Lip}((f-L) \mid U)$ and $C = \max(\mathrm{Lip}(f \mid U), \mathrm{Lip}(f^{-1} \mid U))$. Note that $\kappa(\gamma, U) < 1$ if the number γ and the neighborhood U are sufficiently small. The main result can now be stated more precisely as follows.

Theorem 3. Let $0 < \gamma < 1$, and let U be a neighborhood of the fixed point 0, such that $\kappa(\gamma, U) < 1$. Then for a pair $(\mathscr{F}^s, \mathscr{F}^u)$ of γ-skew foliations on U and a point $p \in U$ the operator $T: \Sigma_p \to \Sigma_p$, defined by (T_1) above, is a contraction with contractive factor $\kappa(\gamma, U)$. The unique fixed point of T is the orbit of $W^s \cap \mathscr{F}_p^u$.

Using e.g. the pair of linear foliations $(\mathscr{E}^s, \mathscr{E}^u)$, we obtain an efficient algorithm, that is easy to implement. This will be discussed in the next section. A proof of Theorem 3 is given in the appendix.

4. Implementation and numerical aspects

Numerical parameters

To obtain an efficient implementation of the operator T introduced in the previous section we use the pair of linear foliations $(\mathscr{E}^s, \mathscr{E}^u)$. The computation of $\pi_s(p)$ and $\pi_u(p)$ for a point p requires the computation of eigenvalues and eigenvectors of the linear part L of f in 0, which can be done using standard methods from numerical linear algebra, see e.g. [3, 11].

The *input to our algorithm* is a point p in some small neighborhood of the fixed point 0. The initial sequence should consist of a finite number of points in some small neighborhood of 0, starting on $\mathscr{E}_p^u$, see also Section 2 for a suitable choice of this initial sequence. The *output* consists of the first N points of the forward orbit of the projection of p onto the stable manifold, along leaves of $\mathscr{F}^u$. The number N, as well as δ, a positive number that is the maximal error between the computed sequence and the actual orbit, are also part of the input.

In view of (T_1) the image $\boldsymbol{x}'$ of a finite sequence $\boldsymbol{x} = \{x_n\}_{0 \le n < n_0}$ of length n_0 under the operator T is a finite sequence $\boldsymbol{x}' = \{x'_n\}_{0 \le n < n_0 - 1}$ of length $n_0 - 1$, defined by

$$x'_n = \begin{cases} [p, f^{-1}(x_1)], & \text{if } n = 0, \\ [f(x'_{n-1}), f^{-1}(x_{n+1})], & \text{if } 0 < n < n_0 - 1, \end{cases} \tag{T_2}$$

where $[y, z]$ denotes the point $\mathscr{E}^u_y \cap \mathscr{E}^s_z = \pi_s(y) + \pi_u(z)$, for y, z near $0 \in \mathbb{R}^d$. We now determine the minimal number M of interations of T needed to guarantee that the error, viz the distance between the computed sequence and the fixed point of T, does not exceed the a priori error bound δ. A first estimate is obtained as follows. If r is the size of the neighborhood of 0 on which the algorithm converges, and $\kappa_0 := \max(\lambda_s, \lambda_u^{-1})$ is taken as an approximation of the contractive factor $\kappa(\gamma, U)$ of the operator T, we see that after M iterations the distance of the current sequence to the fixed point is at most $\kappa_0^M r$. Therefore an estimated upper bound for the number of iterations is

$$M = \frac{\log(\delta/r)}{\log \kappa_0}. \tag{2}$$

Hence, if the output sequence is to be of length N, the initial sequence should consist of at least $M + N$ points. Recall that λ_s (λ_u) is the largest (smallest) modulus of the contracting (expanding) eigenvalues of L. They can be computed using standard methods from numerical linear algebra, again see [3].

We shall see below that this value of M is a very pessimistic estimate for the number of iterations that guarantees a certain maximal error bound. A better approach is to determine the maximal distance δ_0 between successive sequences in order to achieve the error bound δ. If ξ is the fixed point of T, then

$$\begin{aligned}\sigma(T^n \boldsymbol{x}, \xi) &\leq \sigma(T^n \boldsymbol{x}, T^{n+1} \boldsymbol{x}) + \sigma(T^{n+1} \boldsymbol{x}, \xi) \\ &\leq \delta_0 + \kappa_0 \sigma(T^n \boldsymbol{x}, \xi).\end{aligned}$$

Hence

$$\sigma(T^n \boldsymbol{x}, \xi) \leq \frac{\delta_0}{1 - \kappa_0}.$$

Therefore our iteration may stop as soon as the distance between two successive sequences is smaller than $\delta_0 := (1 - \kappa_0)\delta$. This is the approach of Example 4 below. The number of iterations, determined according to the latter criterion, turns out to be much smaller than the upper bound M given above. The latter number is used to determine the length of the initial sequence. These observations are applied in the next two examples.

Performance of the algorithm

Example 4. We consider the map $f: \mathbb{R}^3 \to \mathbb{R}^3$ defined by $f = \Phi \circ L \circ \Phi^{-1}$, where L is the linear map defined by $L(x, y, z) = (\lambda_x x, \lambda_y y, \lambda_z z)$, and Φ is defined by $\Phi(x, y, z) = (x, y, z - ax^2 - by^2)$. We consider various cases in

Computed orbit	Exact orbit
−1.200014	−1.200000
−0.700011	−0.700000
−0.392929	−0.392920
−0.206861	−0.206854
−0.096254	−0.096248
−0.032332	−0.032327
0.003016	0.003020
0.021135	0.021138
0.029082	0.029084
0.031212	0.031214

Figure 3
Comparing a computed orbit in the stable manifold with the exact orbit. Parameter values: $p = (2, 1, 0) \in E^s$, $\lambda_x = 0.8$, $\lambda_y = 0.9$, $\lambda_z = 1.2$, $a = 0.4$, $b = -0.4$. The error bound δ is equal to 10^{-4}. Hence $M = 87$, whereas the actual number of iterations turns out to be 26.

which $0 < \lambda_x, \lambda_y < 1$ and $\lambda_z > 1$, so E^s is the (x, y)-plane, and E^u is the z-axis. Furthermore, $W^s = \Phi(E^s \times \{0\})$ is the graph of the function $(x, y) \mapsto z = ax^2 - by^2$. Therefore we can compare the numerical results of our algorithm with the *exact* values. In Fig. 3 we compare the output of our algorithm for the initial point $p = (2, 1, 0)$ with the exact orbit of $z_0 = \mathcal{O}^u_{x_0} \cap W^s$, viz $z_0 = (2, 1, -1.2)$. In this special example corresponding points in the computed and exact sequences have the same x- and y-coordinates, so we merely need to compare their z-coordinates. The contractive factor is approximately 0.9. After 26 iterations the estimated error between two successive sequences is less than $(1 - 0.9)10^{-4}$, so the distance between the computed sequence and the fixed point is smaller than 10^{-4}. The a priori upper bound on the number of iterations is $M = 87$, see (2). As said before, the latter value of M is used to determine the length of the initial sequence of points.

Running the algorithm for several contractive factors reveals that the actual number of iterations is much smaller than the theoretical upper bound M. For $\kappa_0 = 0.95$ and $\kappa_0 = 0.99$ we have $M = 180$ and $M = 916$, respectively. The actual number of iterations in these cases is 36 and 57, respectively. In Fig. 4 we consider various values of λ_s and λ_u^{-1}, and compare the predicted rate of convergence, viz $\kappa_0 = \max(\lambda_s, \lambda_u^{-1})$, and the actual rate of convergence, viz

$$\sigma_0 := \frac{\sigma(T^{n+1}x, T^n x)}{\sigma(T^n x, T^{n-1}x)},$$

for large values of n. Also in this case our estimate turns out to be too pessimistic, especially if contractive factors are small.

Example 5: A Poincaré map. The algorithm is applied to study the Poincaré, or first-return, map of the parametrically forced 2π-periodic system

$$\ddot{x} + (\alpha^2 + \beta p(t))V'(x) = 0,$$

κ_0	σ_0
0.6	0.1439
0.7	0.248
0.8	0.2898
0.9	0.6753
0.95	0.7521

Figure 4
Convergence rate in theory (left column) and practice (right column).

where $p(t) = \cos t$ and $V(x) = \frac{1}{2}x^2 + \frac{1}{4}x^4$. This is a slightly generalized version of the parametrically forced pendulum, cf. [1]. We consider the central 2π-periodic solution $x = \dot{x} = 0$. To this end we study the Poincaré map of the equivalent first order system

$$\begin{cases} \dot{x} = y \\ \dot{y} = -(\alpha^2 + \beta p(t))V'(x) \\ \dot{t} = 1 \end{cases}$$

defined on the generalized phase space with coordinates $(x, y, t) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}/(2\pi\mathbb{Z})$, where $y = \dot{x}$. The Poincaré map and its inverse are obtained by numerical integration of the system during time 2π and -2π, respectively. The algorithm computes an approximation of the local invariant manifolds of the fixed point $x = y = 0$, for $\alpha = 0.5$ and $\beta = 0.2$. For these values of the parameters the solution $x = y = 0$ is a hyperbolic fixed point of P, see e.g. [1].

The linearized Poincaré map is used to obtain the tangent vectors at $0 \in \mathbb{R}^2$ to the stable and unstable manifolds. The lines parallel to these vectors serve as leaves for the *skew* foliations $\mathscr{F}^s$ and $\mathscr{F}^u$, cf. Section 3. These local manifolds are extended to global approximations via iteration of the Poincaré-map, see Fig. 5. A rigorous proof of the existence of transversal homoclinic intersections of the *global* invariant manifolds of the fixed point is beyond the scope of this paper, although the numerical results provide some evidence for this.

Numerical issues and optimizations

Computation of the inverse map. In many cases using the inverse of f causes no computational problems, since either an explicit expression for $f^{-1}(x_{n+1})$ is known, or this value is easily calculated numerically, as e.g. in the case where f is a Poincaré-map, associated with a periodic trajectory of a vector field, see Example 5. Note that in the latter case the linear part of f can be obtained by numerically solving the variational equation along this periodic trajectory. Also f and f^{-1} are obtained by numerical integration, in forward or backward direction, respectively.

In some applications it may be inefficient to compute f^{-1}, or even impossible, as in some infinite dimensional contexts. In these cases it is

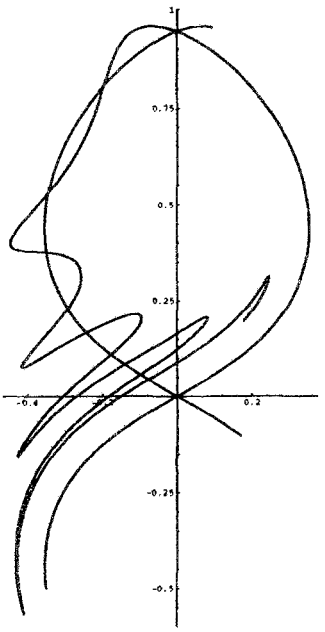

Figure 5
The invariant manifolds of the fixed point of the Poincaré map.

possible to use an alternative definition of the contraction T. Putting $\Delta := x_{n+1} - f(x_n)$, we see that $f^{-1}(x_{n+1}) = x_n + Df^{-1}(f(x_n))\Delta + O(\|\Delta\|^2)$. Approximating the right hand side of the last equality by $x_n + L^{-1}(x_{n+1} - f(x_n))$, we get a slightly different version of the operator T by replacing $\pi_u(f^{-1}(x_{n+1}))$ in (T_2), for $n \geq 0$, with $\pi_u(x_n) + (L^u)^{-1}\pi_u(x_{n+1} - f(x_n))$. (In fact the approximation used to compute $f^{-1}(x_{n+1})$ is just the first step in a modified Newton's method.)

Continuous systems. The invariant manifold of a hyperbolic singular point p of a vector field, viz a system of autonomous ordinary differential equations, coincide with the invariant manifolds of the time-1 map f of this system. In this case $f(x)$ and $f^{-1}(x)$ are obtained by forward and backward numerical integration during 1 unit of time, starting at the point x. The linear part of this map at p is the image of the linear part of the system at p under the exponential map, which can be computed using methods from numerical linear algebra, see again [3].

Region of convergence. Using the expression for the contractive factor $\kappa(\gamma, U)$, see (1), in many cases we can give an estimate of the size of the neighborhood of the fixed point on which our algorithm converges. Roughly speaking, for fixed γ, the size of this region is inversely proportional to the strength of the nonlinear terms. We have performed some numerical experiments that support this observation. In practical situations our theoretical estimates again turned out to be rather pessimistic, especially if the contraction is strong.

Robustness. If the map f is perturbed slightly, the pair of foliations used in the computation is still γ-skew for the perturbed map. Therefore our method is persistent. It can also be used in continuation-like methods, e.g. to compute the invariant manifolds of a map depending on some continuously changing external parameter. Again this is a common approach in perturbation theory.

5. Pseudohyperbolic fixed points

The algorithm can be extended to compute the strong stable and unstable manifolds of pseudohyperbolic fixed points. We say that the fixed point $0 \in \mathbb{R}^d$ is ϱ-pseudohyperbolic for some $\varrho > 0$, if no eigenvalue of $L := Df(0)$ has modululs ϱ. As in the hyperbolic case there is a linear L-invariant decomposition $\mathbb{R}^d = E_1 \oplus E_2$, such that the eigenvalues of $L_1 := L \mid E_1$ are inside the circle in the complex plane with radius ϱ, and the eigenvalues of $L_2 := L \mid E_2$ are outside this circle. Note that there is a norm on $\mathbb{R}^d$ such that $\lambda_1 := \|L_1\| < \varrho$, and $\lambda_2^{-1} := \|(L_2)^{-1}\| > \varrho^{-1}$.

If $\varrho \leq 1$ the *strong stable* manifold W^{ss} of the map f at 0 is the set of points p that are ϱ-forward asymptotic to 0, i.e. $\lim_{n \to \infty} \varrho^{-n} d(f^n(p), 0) = 0$. This set is a unique, f-invariant, C^∞ manifold, whose tangent space at the fixed point is E_1, see [5], Section 5. If $\varrho = 1$, i.e. if 0 is a hyperbolic fixed point, the strong stable manifold is just the stable manifold as introduced in Section 2. Also here we use the equivalent definition of W^{ss} as the set of points p near 0 for which $\varrho^{-n} f^n(p)$ is uniformly bounded (see e.g. [5]).

If $\varrho \geq 1$ we similarly have a uniquely determined, f-invariant, C^∞ manifold W^{uu}, consisting of those points that are locally ϱ-backward asymptotic to 0. The tangent space of this *strong unstable* manifold at the fixed point is E_2.

Remark 6. If $\varrho < 1$ there is also an f-invariant manifold tangent to E_2 at the fixed point. It is not C^∞ in general: it is of class C^r, where r is an integer such that f is C^r and $\lambda_1 \lambda_2^{-j} < 1$ for $1 \leq j \leq r$. Furthermore, it is generally not unique. This causes algorithmic problems that lie outside the scope of this paper. We are currently developing algorithms to compute normally hyperbolic invariant manifolds. The computation of an f-invariant manifold tangent to E_2 falls naturally within this context.

Henceforth we only deal with the case $\varrho < 1$. Let $\pi_1 \colon \mathbb{R}^d \to E_1$ and $\pi_2 \colon \mathbb{R}^d \to E_2$ be the canonical projections corresponding to the L-invariant splitting $\mathbb{R}^d = E_1 \oplus E_2$.

The orbit of a point in W^{ss} belongs to the space of sequences $\{x_n\}_{n \geq 0}$ for which $\varrho^{-n} x_n$ is uniformly bounded. For p near 0 we consider the space

of sequences whose initial points lie on the leaf through p of some suitably chosen foliation. However, here we are facing the problem that, in general, there is no f-invariant foliation whose leaf through 0 is tangent to E_2. This is related to the fact that, in general, there is no f-invariant cone-field whose axes are parallel to E_2 (like the unstable cone-field in the hyperbolic case). Therefore, we consider a splitting (not necessarily f-invariant) $\mathbb{R}^d = E_1' \oplus E_2'$, such that, for some $\gamma > 0$ and points x, y with $\pi_1'(x) = \pi_1'(y)$,

$$\|\pi_1(x) - \pi_1(y)\| \leq \gamma \|\pi_2(x) - \pi_2(y)\|.$$

In this case we say that the foliation $\mathscr{E}_2'$, whose leaves are parallel to E_2', is γ-near $\mathscr{E}_2$, whose leaves are parallel to E_2. One similarly defines the foliation $\mathscr{E}_1'$, that is γ-near $\mathscr{E}_1'$. The pair of foliations $(\mathscr{E}_1', \mathscr{E}_2')$ is completely similar to the γ-skew foliations, introduced in Section 3.

For a point p near 0, we introduce the space Σ_p, consisting of all sequences $\{x_n\}_{n \geq 0}$, whose initial point lies on the same leaf of $\mathscr{E}_2'$ as p, such that $\varrho^{-n} x_n$ is uniformly bounded. It is not hard to check that Σ_p, endowed with the sup-norm, is a Banach space, cf. Section 3. Let Σ_p' be the space of *bounded sequences*, whose initial points lie on the same leaf of $\mathscr{E}_2'$ as p. We endow this space with a norm by imposing the condition that the linear map $A: \Sigma_p' \to \Sigma_p$, defined by $(A\boldsymbol{x})_n = \varrho^{-n} x_n$, for $\boldsymbol{x} = \{x_i\}_{i \geq 0}$ and $n \geq 0$, is an isometry. Then the operator $T': \Sigma_p' \to \Sigma_p'$, defined by (T_2) in Section 4, is a contraction. Equivalently we have:

Theorem 7. The operator $T: \Sigma_p \to \Sigma_p$, defined by

$$T = A \circ T' \circ A^{-1},$$

is a contraction. Its fixed point is the f-orbit of the point of intersection of W^{ss} and the leaf of $\mathscr{E}_2'$ through p.

In the appendix, where we prove the main theorem, we also indicate how to prove Theorem 7. The algorithm, suggested by Theorem 7, is almost identical to the algorithm discussed in Section 4 (also cf. [6] for similar ideas). Therefore, we omit further details here.

6. Appendix: Proof of the main theorem

We first present some auxiliary results, to be used in the proof of our main result. Let $(\mathscr{F}^s, \mathscr{F}^u)$ be a pair of γ-skew foliations. Leaves of $\mathscr{F}^u$ are

nearly parallel to E^u. More precisely:

Lemma 8. Let $0 < \gamma < 1$. Then there is a neighborhood U of 0 such that for $x, \bar{x}, y, \bar{y} \in U$ satisfying $\mathscr{F}_x^u = \mathscr{F}_{\bar{x}}^u$, and $\mathscr{F}_y^u = \mathscr{F}_{\bar{y}}^u$ we have

$$\|\pi_s(x) - \pi_s(y)\| \leq \gamma \|\pi_u(x) - \pi_u(y)\| + (1+\gamma)\|\pi_s(\bar{x}) - \pi_s(\bar{y})\| + 2\gamma \|\pi_u(\bar{x}) - \pi_u(\bar{y})\|.$$

A similar inequality holds if we interchange u and s.

Proof. *Special case*: $\pi_u(x) = \pi_u(y)$ and $\pi_u(\bar{x}) = \pi_u(\bar{y})$.

Since $\mathscr{F}^u$ is a foliation of class C^1, its leaves are level sets of a C^1-function $H: E^s \times E^u \to E^s$. We may assume that $H(\xi, 0) = \xi$, for $\xi \in E^s$, provided ξ is sufficiently near 0. Let $(\xi, 0)$ and $(\bar{\xi}, 0)$ be the points of intersection of $E^s \times \{0\}$ with the leaves $\mathscr{F}_x^u$ and $\mathscr{F}_y^u$, respectively. Then $H(x) = H(\bar{x}) = \xi$, and $H(y) = H(\bar{y}) = \bar{\xi}$. Therefore

$$\begin{aligned} \xi - \bar{\xi} &= H(\pi_s(x), \pi_u(x)) - H(\pi_s(y), \pi_u(x)) \\ &= D_s H(\xi_0, \pi_u(x))(\pi_s(x) - \pi_s(y)), \end{aligned}$$

where the last equality follows from the mean value theorem for some ξ_0 on the line segment E^s between $\pi_s(x)$ and $\pi_s(y)$. Here $D_s H(\xi_0, \pi_u(x))$ is the derivative of the map $H(\cdot, \pi_u(x)): E^s \to E^s$. Since $D_s H(\xi_0, 0)$ is the identity map on E^s, we have

$$1 - \frac{1}{3}\gamma \leq \|D_s H(\xi_0, \pi_u(x))\| \leq 1 + \frac{1}{3}\gamma,$$

provided $(\xi_0, \pi_u(x))$ is near 0. The latter condition holds if x, y, $\bar{x}$ and $\bar{y}$ are in some small neighborhood of 0. Under these conditions we have

$$\left(1 - \frac{1}{3}\gamma\right)\|\pi_s(x) - \pi_s(y)\| \leq \|\xi - \bar{\xi}\| \leq \left(1 + \frac{1}{3}\gamma\right)\|\pi_s(x) - \pi_s(y)\|.$$

A similar inequality holds if we replace x with $\bar{x}$, and y with $\bar{y}$. Combining these inequalities and using the fact that $0 < \gamma < 1$ we obtain

$$\|\pi_s(x) - \pi_s(y)\| \leq (1+\gamma)\|\pi_s(\bar{x}) - \pi_s(\bar{y})\|,$$

which completes the proof of this special case.

The general case. Let $p = \mathscr{E}_x^s \cap \mathscr{F}_y^u$ and $q = \mathscr{E}_{\bar{x}}^s \cap \mathscr{F}_{\bar{y}}^u$. Since $\mathscr{F}^u$ is γ-skew we have

$$\begin{aligned} \|\pi_s(p) - \pi_s(y)\| &\leq \gamma \|\pi_u(p) - \pi_u(y)\| \\ &= \gamma \|\pi_u(x) - \pi_u(y)\|. \end{aligned}$$

In view of the first part of the proof we have

$$\|\pi_s(p) - \pi_s(x)\| \leq (1+\gamma)\|\pi_s(q) - \pi_s(y)\|.$$

Using these results we derive:

$$\begin{aligned}\|\pi_s(x) - \pi_s(y) &\leq \|\pi_s(x) - \pi_s(p)\| + \|\pi_s(p) - \pi_s(y)\| \\ &\leq (1+\gamma)\|\pi_s(q) - \pi_s(\bar{x})\| + \gamma\|\pi_u(x) - \pi_u(y)\| \\ &\leq (1+\gamma)\|\pi_s(\bar{y}) - \pi_s(\bar{x})\| + (1+\gamma)\|\pi_s(q) - \pi_s(\bar{y})\| \\ &\quad + \gamma\|\pi_u(x) - \pi_u(y)\| \\ &\leq (1+\gamma)\|\pi_s(\bar{y}) - \pi_s(\bar{x})\| + (1+\gamma)\gamma\|\pi_u(\bar{x}) - \pi_u(\bar{y})\| \\ &\quad + \gamma\|\pi_u(x) - \pi_u(y)\|.\end{aligned}$$

Since $0 < \gamma < 1$ the result follows. □

Lemma 9. Let $\varepsilon > 0$. There is a neighborhood U of 0 such that for x, $y \in U$ we have:

$$\|\pi_s f(x) - \pi_s f(y)\| \leq (\lambda_s + \varepsilon)\, d(x, y),$$
$$\|\pi_u f^{-1}(x) - \pi_u f^{-1}(y)\| \leq (\lambda_u^{-1} + \varepsilon)\, d(x, y).$$

Proof. We only prove the first inequality. The proof of the second one is similar.

$$\begin{aligned}\|\pi_s f(x) - \pi_s f(y)\| &\leq \|\pi_s L(x-y) + \pi_s(f-L)x - \pi_s(f-L)y\| \\ &\leq \|\pi_s L(x-y)\| + \mathrm{Lip}(\pi_s(f-L))\, d(x, y) \\ &\leq \lambda_s\, d(x, y) + \varepsilon\, d(x, y).\end{aligned}$$

□

Remark 10. An immediate consequence of this lemma is the f-invariance of the cone-fields $\{C_x^s(\gamma)\}$ and $\{C_x^u(\gamma)\}$, for $0 < \gamma < 1$ and x ranging over some neighborhood U_γ of 0, in the sense described in Section 3.2. More precisely, the neighborhood U_γ should satisfy $\mathrm{Lip}(\pi_s(f-L) \mid U_\gamma) < \gamma\varepsilon$. It is not difficult to prove that, for x, $y \in U_\gamma$, with $y \in C_x(\gamma)$:

$$\|\pi_s f(x) - \pi_s f(y)\| \leq \gamma(\lambda_s + \varepsilon)(\lambda_u^{-1} + \varepsilon)\|\pi_u f(x) - \pi_u f(y)\|.$$

Therefore $f(y) \in C_{f(x)}(\gamma)$.

Remark 11. If 0 is a ϱ-pseudohyperbolic fixed point of f it is not hard to show that, in the notation of Section 5:

$$\|\pi_1 f(\varrho^{-1}x) - \pi_1 f(\varrho^{-1})y)\| \leq (\lambda_1 \varrho^{-1} + \varepsilon)\, d(x, y),$$
$$\|\pi_2 f^{-1}(\varrho x) - \pi_2 f^{-1}(\varrho y)\| \leq (\varrho\lambda_2^{-1} + \varepsilon)\, d(x, y).$$

Proof of the main theorem: Let $\boldsymbol{x} = \{x_n\}_{n \geq 0}$, $\boldsymbol{y} = \{y_n\}_{n \geq 0} \in \Sigma_p$ be two sequences, and let $\boldsymbol{x}' = \{x'_n\}_{n \geq 0}$, $\boldsymbol{y}' = \{y'_n\}_{n \geq 0} \in \Sigma_p$ be their images under T. We shall prove inductively that

$$d(x'_m, y'_m) \leq \kappa(\gamma, U)\sigma(\boldsymbol{x}, \boldsymbol{y}), \tag{3}$$

for $m = 0, 1, \ldots$. The proof of the base case, viz $m = 0$, is a simple version of the proof for the general case. So assume that (3) has been proven for $0 \leq m \leq n-1$.

Since $\mathscr{F}^u_{x'_n} = \mathscr{F}^u_{f(x'_{n-1})}$ and $\mathscr{F}^u_{y'_n} = \mathscr{F}^u_{f(y'_{n-1})}$ we have, in view of Lemma 8:

$$\begin{aligned}\|\pi_s(x'_n) - \pi_s(y'_n)\| \leq{}& (1+\gamma)\|\pi_s f(x'_{n-1}) - \pi_s f(y'_{n-1})\| \\ &+ 2\gamma\|\pi_u f(x'_{n-1}) - \pi_u f(y'_{n-1})\| \\ &+ \gamma\|\pi_u(x'_n) - \pi_u(y'_n)\|.\end{aligned}$$

Let C_1 be the Lipschitz constant of $f \mid U$, then we get, using Lemma 9:

$$\begin{aligned}\|\pi_s(x'_n) - \pi_s(y'_n)\| \leq{}& (1+\gamma)(\lambda_s + \varepsilon)\, d(x'_{n-1}, y'_{n-1}) \\ &+ 2\gamma C_1\, d(x'_{n-1}, y'_{n-1}) \\ &+ \gamma\|\pi_u(x'_n) - \pi_u(y'_n)\|.\end{aligned}$$

Using the induction hypothesis we get:

$$\begin{aligned}\|\pi_s(x'_n) - \pi_s(y'_n)\| \leq{}& ((1+\gamma)(\lambda_s + \varepsilon) + 2\gamma C_1)\sigma(\boldsymbol{x}, \boldsymbol{y}) \\ &+ \gamma\|\pi_u(x'_n) - \pi_u(y'_n)\|.\end{aligned} \tag{4}$$

Arguing similarly we get, using $\mathscr{F}^s_{x'_n} = \mathscr{F}^s_{f^{-1}(x_{n+1})}$ and $\mathscr{F}^s_{y'_n} = \mathscr{F}^s_{f^{-1}(y_{n+1})}$:

$$\begin{aligned}\|\pi_u(x'_n) - \pi_u(y'_n)\| \leq{}& ((1+\gamma)(\lambda_u^{-1} + \varepsilon) + 2\gamma C_2)\sigma(\boldsymbol{x}, \boldsymbol{y}) \\ &+ \gamma\|\pi_s(x'_n) - \pi_s(y'_n)\|,\end{aligned} \tag{5}$$

where C_2 is the Lipschitz constant of $f^{-1} \mid U$. Since $C = \max(C_1, C_2)$, we conclude that

$$d(x'_n, y'_n) \leq (1-\gamma)^{-1}((1+\gamma)(\max(\lambda_s, \lambda_u^{-1}) + \varepsilon) + 2\gamma C)\sigma(\boldsymbol{x}, \boldsymbol{y}).$$

To see this, we use (4) if $\|\pi_u(x'_n) - \pi_u(y'_n)\| \leq \|\pi_s(x'_n) - \pi_s(y'_n)\|$, and (5) otherwise. This completes the proof of our main result. □

Acknowledgement

We are grateful to H. W. Broer and F. Takens for reading the draft version and giving helpful suggestions, and to F. Wubs for valuable remarks on the numerical aspects of our method. We also thank the anonymous referees for their remarks, that improved the presentation of the algorithm.

References

[1] H. W. Broer and G. Vegter, *Bifurcational Aspects of Parametric Resonance*, vol. 1 of *Dynamics Reported (New Series)*, chapt. 1, pp. 1–53. Springer Verlag, Berlin–Heidelberg–New York 1992.
[2] R. L. Devaney, *An Introduction to Chaotic Dynamical Systems*. Addison-Wesley Publ Co., Redwood City, CA 1987.
[3] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Johns Hopkins University Press, Baltimore 1983.
[4] J. Hadamard, *Sur l'itération et les solutions asymptotiques des equations différentielles*. Bull. Soc. Math. France, *29*, 224–228 (1901).
[5] M. W. Hirsch, C. C. Pugh, and M. Shub, *Invariant Manifolds*. Springer-Verlag, Berlin 1977.
[6] A. J. Homburg, *On the Computation of Hyperbolic Sets and their Invariant Manifolds*. Technical Report 68, Institut für Angewandte Analysis und Stochastik, Berlin 1993.
[7] F. Ma and T. Küpper, *Numerical calculation of invariant manifolds for maps*. J. Num. Linear Alg. Appl., *1,2*, 141–150 (1994).
[8] H. E. Nusse and J. A. Yorke, *A procedure for finding numerical trajectories on chaotic saddles*. Physica D, *36*, 137–156 (1989).
[9] J. Palis and W. de Melo, *Geometric Theory of Dynamical Systems*. Springer-Verlag, New York 1982.
[10] J. Palis and F. Takens, *Hyperbolicity & sensitive chaotic dynamics at homoclinic bifurcations*, vol. 35 of Cambridge studies in advanced math. Cambridge University Press 1993.
[11] T. S. Parker and L. O. Chua, *Practical Numerical Algorithms for Chaotic Systems*. Springer-Verlag, Berlin 1989.
[12] O. Perron, *Ueber Stabilität und Asymptotisches Verhalten der Lösungen eines Systemes endlicher Differenzengleichungen*. J. Reine Angew. Math., *161*, 41–46 (1929).
[13] M. Shub, *Global Stability of Dynamical Systems*. Springer-Verlag, Berlin 1977.
[14] S. van Gils and A. Vanderbauwhede, *Center manifolds and contractions on a scale of Banach spaces*. J. Funct. Analysis, pp 209–224 (1987).
[15] A. Vanderbauwhede, *Center Manifolds, Normal Forms and Elementary Bifurcations*. In *Dynamics Reported*, vol. 2, pp 89–170. John Wiley & Sons Ltd and B.G. Teubner, Stuttgart 1989.
[16] H. E. S. Westerveld, Numerieke Bepaling van Invariante Variëteiten (in Dutch). Master's thesis, University of Twente 1990.

Abstract

We present a method for the numerical computation of invariant manifolds of hyperbolic and pseudohyperbolic fixed points of diffeomorphisms. The derivation of this algorithm is based on well-known properties of (almost) invariant foliations. Numerical results illustrate the performance of our method.

(Received: April 12, 1994; revised: September 10, 1994)