Krauskopf, B., & Osinga, HM. (2003). *Computing geodesic level sets on global (un)stable manifolds of vector fields*.
<https://doi.org/10.1137/030600180>

Early version, also known as pre-print

Link to published version (if available):
[10.1137/030600180](https://doi.org/10.1137/030600180)

[Link to publication record in Explore Bristol Research](#)
PDF-document

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

COMPUTING GEODESIC LEVEL SETS ON GLOBAL (UN)STABLE MANIFOLDS OF VECTOR FIELDS

BERND KRAUSKOPF* AND HINKE M. OSINGA*

7th May 2003

Abstract.

Many applications give rise to dynamical systems in the form of a vector field with a phase space of moderate dimension. Examples are the Lorenz equations, mechanical and other oscillators, and models of spiking neurons. The key to understanding the global dynamics of such a system are the stable and unstable manifolds of the saddle points, of the saddle periodic orbits and, more generally, of all invariant manifolds of saddle-type. Except in very special circumstances the (un)stable manifolds are global objects that cannot be found analytically but need to be computed numerically. This is a nontrivial task when the dimension of the manifold is larger than one.

In this paper we present a general algorithm to compute the k -dimensional unstable manifold of an equilibrium or periodic orbit (or more general normally hyperbolic invariant manifold) of a vector field with an n -dimensional phase space, where $1 < k < n$. Stable manifolds are computed by considering the flow for negative time. The key idea is to view the unstable manifold as a purely geometric object, hence disregarding the dynamics on the manifold, and compute it as a list of approximate geodesic level sets, which are (topological) $(k - 1)$ -spheres. Starting from a $(k - 1)$ -sphere in the linear eigenspace of the equilibrium or periodic orbit, the next geodesic level set is found in a local (and changing) coordinate system given by hyperplanes perpendicular to the last geodesic level set. In this setup the mesh points defining the approximation of the next geodesic level set can be found by solving a boundary value problem. By appropriately adding or removing mesh points it is ensured that the mesh that represents the computed manifold is of a prescribed quality.

The general algorithm is presently implemented to compute two-dimensional manifolds in a phase space of arbitrary dimension. In this case the geodesic level sets are topological circles and the manifold is represented as a list of ribbons between consecutive level sets. We use color to distinguish between consecutive ribbons or to indicate geodesic distance from the equilibrium or periodic orbit, and we also show how geodesic level sets change with increasing geodesic distance. This is very helpful when one wants to understand the often very complicated embeddings of two-dimensional (un)stable manifolds in phase space.

The properties and performance of our method are illustrated with several examples, including the stable manifold of the origin of the Lorenz system, a two-dimensional stable manifold in a four-dimensional phase space arising in a problem in optimal control, and a stable manifold of a periodic orbit that is a Möbius strip. Each illustration is accompanied by an animation (supplied with this paper).

Key words. Global stable and unstable manifolds, numerical method, geodesic distance

AMS subject classifications. 37C10, 34K19, 37M20

1. Introduction. Many dynamical systems arising in applications can be written as a system of ordinary differential equations. Well-known examples are the Lorenz system, Chua's circuit and any number of periodically forced systems arising in applications. See for example, the text books [8, 24, 30] and further references therein, and also the examples in Section 4. In order to understand the global dynamics of such a system it is necessary to compute its equilibria, periodic orbits, and possibly other normally hyperbolic invariant sets, such as invariant tori. If these objects are of saddle-type one also needs to know their stable and unstable manifolds. These global manifolds organize the dynamics: stable manifolds (of codimension one) may form boundaries of basins of attraction and intersections of stable and unstable manifolds are responsible for complicated dynamics and chaos. These important objects can generally be found by numerical methods only.

*Department of Engineering Mathematics, University of Bristol, Bristol BS8 1TR, UK.

To fix notation we consider a vector field in its general form

$$\dot{x} = f(x), \quad (1.1)$$

where x is from an n -dimensional phase space, which we take to be $\mathbb{R}^n$ for simplicity, and $f : \mathbb{R}^n \mapsto \mathbb{R}^n$ is sufficiently smooth. This means in particular that a stable manifold of (1.1) can be computed as an unstable manifold by reversing time. Note that we consider all possible parameters of the vector field to be constant. In order to see what happens to the manifolds under variation of parameters one will have to compute them for individual fixed values of the parameters.

To keep this introduction simple, we explain the basic idea with the example of an unstable manifold of a saddle point $x_0 \in \mathbb{R}^n$. Suppose that x_0 has $1 < k < n$ eigenvalues with positive real parts and $(n - k)$ eigenvalues with negative real parts (counted with multiplicity). The Unstable Invariant Manifold Theorem (see, for example, [24, 27]) guarantees that in a neighborhood of x_0 there exists the local unstable manifold $W_{\text{loc}}^u(x_0)$, which is tangent to the unstable (generalized) eigenspace $E^u(x_0)$ of the Jacobian $Df(x_0)$ of (1.1) at x_0 . Furthermore, $W_{\text{loc}}^u(x_0)$ is as smooth as f . The local unstable manifold $W_{\text{loc}}^u(x_0)$ can be globalised by applying the flow ϕ^t of (1.1) to obtain the k -dimensional (immersed) unstable manifold

$$\begin{aligned} W^u(x_0) &= \{x \in \mathbb{R}^n \mid \lim_{t \rightarrow -\infty} \phi^t(x) = x_0\} \\ &= \lim_{t \rightarrow \infty} \phi^t(W_{\text{loc}}^u(x_0)). \end{aligned}$$

This automatically implies that $W^u(x_0)$ is also as smooth as f .

It is a natural idea to start close to x_0 with a small topological $(k - 1)$ -sphere S_δ (a topological circle for $k = 2$) approximately in $W_{\text{loc}}^u(x_0)$ at distance δ from x_0 , and then ‘grow’ this sphere to obtain $W^u(x_0)$ further away from x_0 . In practice, we take $S_\delta \in E^u(x_0)$, but one could also start with a higher-order approximation of $W^u(x_0)$ close to x_0 .

In the special case $k = 1$ of computing a one-dimensional manifold, the initial $(k - 1)$ -sphere S_δ consists of two points (approximately) on $W_{\text{loc}}^u(x_0)$ at distance δ from x_0 , one on each side of x_0 . These two points can be grown efficiently by integration, that is, by evolving them with the flow ϕ^t of (1.1). In other words, computing one-dimensional unstable manifolds simply amounts to integrating $\dot{x}$ from two suitable initial conditions.

The situation becomes much more complicated for $k \geq 2$. One may be tempted to simply evolve S_δ under the flow ϕ^t of (1.1) also in this case. The idea behind this is that $W^u(x_0)$ can be represented as the family $\{\phi^t(S_\delta)\}_{t \in \mathbb{R}}$ of all t -images of S_δ , parametrized by the (integration) time t . However, this representation does not produce a good numerical approximation of $W^u(x_0)$, because S_δ deforms very rapidly under the flow ϕ^t due to the dynamics on $W^u(x_0)$. For example, it stretches out along the strong unstable direction (if present) and develops very large aspect ratios. Even when all eigenvalues have about the same strength, the initial circle S_δ generally gets much distorted further away from x_0 due to nonlinear effects; see the example in Section 4.4. As a consequence, any initial mesh representing S_δ generally deteriorates so rapidly that simply evolving it under the flow does not result in a good mesh representation of $W^u(x_0)$.

Computing higher-dimensional stable and unstable manifolds of vector fields is quite challenging and an active field of research. A number of algorithms have been devised for this task [4, 5, 7, 9, 10, 13, 17, 18, 20, 22, 25, 26]. The method in [4, 5] is

special as it computes a box covering of $W^u(x_0)$; it has been implemented for arbitrary n and is independent of k . All other algorithms are implemented specifically for the case of two-dimensional manifolds and produce a triangulation of the manifold that is built up by starting near the saddle point. These methods differ in the way they deal with the problem of mesh deterioration; see Section 2 for a more detailed review.

In this paper we present a general algorithm for computing a k -dimensional unstable manifold of a vector field with an n -dimensional phase space. Its formulation in terms of geodesic level sets develops further previous work in [22] for the specific case $k = 2$ and $n = 3$. The method is now fully implemented for the case $k = 2$ and any n . This includes the case of two-dimensional stable and unstable manifolds of periodic orbits. (We remark that an implementation for $k \geq 3$ is work for the future. It would have to deal with serious issues of a suitable simplicial representation and with the problem of visualizing three- or higher-dimensional manifolds.)

The key idea is to step completely away from evolving an existing mesh under the flow. To this end we consider $W^u(x_0)$ purely as a *geometric object*, namely as a family $\{S_\eta\}_{\eta \in \mathbb{R}^+}$ of geodesic level sets parametrized by their *geodesic distance η along the manifold* from x_0 ; see Section 2 for the precise definition. The manifold is grown by computing more and more (approximate) geodesic level sets on it. This leads to a mesh whose quality can be guaranteed irrespective of the dynamics on the manifold, allowing for the statement of convergence in Section 5.

More specifically, we start from an initial level set S_δ in $E^u(x_0)$, and then approximate a sequence of level sets $\{S_{\eta_i}\}_{0 \leq i \leq l}$ until a prespecified geodesic distance D from x_0 is reached. In other words, $\eta_0 = \delta$ and $\eta_l \geq D$. Each level set S_{η_i} is approximated by a $(k - 1)$ -dimensional simplicial complex C_i with mesh points M_i . In the course of the computation the manifold is grown by adding a k -dimensional simplicial representation generated by C_i and the newly computed C_{i+1} . For $k = 2$ each set C_i is a piecewise linear representation of a topological circle, and in each step a band or ribbon is added to the manifold; see already the examples in Section 4.

Given the mesh points in M_i the next mesh points in M_{i+1} can be computed point by point in the local coordinate system given by the (approximate) tangent space and its orthogonal complement $\mathcal{F}_r$ at each mesh point r of M_i . We prescribe that the new mesh point in M_{i+1} lies in $\mathcal{F}_r$ at distance Δ_i from r . Provided Δ_i is not too large, such a point can be found by solving a suitable boundary value problem; for the details see Section 3. The step size Δ_i is adapted according to the curvature of the manifolds. If points move too far from each other within the level set, extra points are added in M_{i+1} . This is always done in such a way that the interpolation only takes place within the prescribed mesh bounds. Similarly, points in M_{i+1} are removed if they come too close to each other.

A main goal of this paper is to illustrate and discuss the performance of our algorithm with several examples, some test examples, and some to highlight its usefulness for the study of systems arising in applications. The interested reader is encouraged to look ahead to Section 4 where we show:

1. the two-dimensional stable manifold of the origin of the Lorenz system for a non-standard choice of parameters when there is a figure-eight shaped periodic orbit as the attractor,
2. the two-dimensional stable manifold in a four-dimensional phase space arising in a problem in optimal control,
3. the two-dimensional stable manifold, in fact a Möbius strip, of a periodic orbit in the ζ^3 model,

4. a constructed example that illustrates a possible geometric limitation of our method, and
5. the two-dimensional stable manifold of the origin of the Lorenz system for the standard choice of parameters, where we use a color coding of the geodesic distance to show how the manifold spirals into the well-known Lorenz attractor.

All examples are accompanied by animations.

The outline of this paper is as follows. We first discuss the conceptual background in Section 2, where we also review other methods for computing global invariant manifolds. We then give a complete description of the general algorithm in Section 3 and discuss some details of the implementation for $k = 2$. The longer Section 4 consists of the detailed exposition of the examples mentioned above. Finally, we discuss the correctness of the general algorithm in Section 5 and draw some conclusions in Section 6.

2. Background and concepts. As was already mentioned in the introduction, there are several possibilities of parametrizing a k -dimensional unstable manifold $W^u(x_0)$ of a saddle point x_0 . Anyone schooled in dynamical systems will know that $W^u(x_0)$ consists of orbits that can be parametrized by their intersection points with a suitable small $(k - 1)$ -sphere S_δ in $W_{\text{loc}}^u(x_0)$. The dual parametrization is given by the t -images of S_δ , that is,

$$W^u(x_0) = \{\phi^t(S_\delta)\}_{t \in \mathbb{R}}, \quad (2.1)$$

which has the advantage that the parameter space is one-dimensional. Because it is given by the dynamics, this parametrization has a number of nice properties. All elements of this parametrization are smooth topological $(k - 1)$ -spheres, just as S_δ . This property holds no matter what the manifold looks like, or whether it converges to an attractor or not. (In the special case $k = 2$ this means that the two-dimensional manifold $W^u(x_0)$ is a one-parameter family of smooth simple closed curves.) However, as discussed in the introduction, the problem is that the smooth topological $(k - 1)$ -spheres of this parametrization generally deform rapidly under the dynamics when t is increased. As a consequence, quickly the topological $(k - 1)$ -spheres are very far from being ‘nice (hyper)spheres’. The reason for this is that the initial sphere does not grow uniformly in all directions under the flow ϕ^t .

This is why simply computing (approximations to) $\phi^t(S_\delta)$ for a discrete set of integration times is completely impractical as a means of obtaining a mesh approximation of $W^u(x_0)$. Indeed, a strategy is needed to counteract this problem when a topological $(k - 1)$ -sphere is evolved under the flow. Appropriate methods have been designed by Guckenheimer and Worfolk [10], Worfolk [34] and Johnson, Jolly and Kevrekidis [17], implemented for the case $k = 2$ and $n = 3$ of growing topological circles to cover a two-dimensional manifold, starting with S_δ . Guckenheimer and Worfolk [10, 34] rescale the vector field so that the tangential component to the last circle is practically zero and the circle is grown in the radial direction by integration. The idea is to move from one approximate geodesic level set to the next with the help of the flow. This requires that the rescaled vector field points radially outward everywhere along the last circle, which is why this approach has serious difficulties in the case of complex conjugate eigenvalues. Also, the first circle must be chosen such that this condition is satisfied, which means that one may need to start with an ellipse rather than S_δ . Johnson *et al.* [17] use a parametrization by arclength of the trajectories. At each step of their computation the mesh points on the furthest circle

are integrated up to a specified arclength, leading to a new circle, on which a uniform mesh is selected by interpolation between the integration points. The lack of control on the interpolation error makes it difficult to assess the accuracy of the computation, a problem that may be even more pronounced for $k > 2$.

A different approach is to view the manifold as a family of orbits. Doedel [7] computes two-dimensional manifolds by following an orbit of a prescribed arclength by continuation with the package AUTO [6], where the angle of the orbit with a reference direction near the equilibrium acts as continuation parameter. The initial condition of each orbit lies again on a small circle $S_\delta \subset E^u(x_0)$. His method is very accurate and is particularly suited for the case $k = 2$ where the manifold can be represented as a one-parameter family of orbits. However, it does seem to be more difficult to generalize this to $k > 2$. Furthermore, the method suffers from an uneven distribution of mesh points: many are needed near the equilibrium to obtain a good mesh near the end points of the orbits.

The method of Henderson [13] is similar in spirit, but tries to better control the mesh by using only the local information of the orbits. His method starts with integrating a particular orbit together with higher-order manifold information. This results in a string of polyhedral patches along the orbit, called a fattened trajectory. Then the new boundary of the computed part of the manifold is found and the next fattened trajectory is added by starting from a suitable point along this boundary. The computation stops when all fattened trajectories have been computed up to a specified arclength from S_δ .

In the method by Guckenheimer and Vladimirovsky [9] the unstable manifold is grown by adding new mesh points locally. Their method is fast by keeping the integration of the vector field to a minimum. It is implemented for the case $k = 2$ and adds a simplex locally to the manifold at each step. All but one point of the simplex are known mesh points, and the position of the new point is found by solving a PDE in an Eulerian framework, which formalises the invariance condition, with an upwind method. Where the next triangle is added is determined in a clever way by the upwinding (essentially by the local direction of the vector field). This allows one to compute the manifold, for example, up to a predetermined arclength of the orbits on the manifold (which are approximated by the mesh structure). Due to the upwinding, the algorithm does not produce fattened trajectories, but tends to add triangles in a way similar to a growth method.

Finally, Dellnitz and Hohmann [4, 5] do not use any specific form of parametrization. They consider the time- τ map of the flow of (1.1) for some fixed $\tau > 0$ and cover $W_{\text{loc}}^u(x_0)$ with n -dimensional boxes. This box covering is then evolved in a prescribed region A to obtain a box covering of the connected piece of $W^u(x_0)$ in A that contains x_0 . This covering method is non-uniform and depends on the dynamics on the manifold. The implementation of this algorithm is independent of the dimension k of the manifold. However, the practical implementation of reliably detecting when the image of one box intersects another box (for example, by using test points) remains a challenge already for $n = 3$.

All above algorithms (and this includes our own) can, in principle, be used for higher-dimensional manifold computations, but present their own specific challenges of increased computational complexity for $k \geq 3$. Furthermore, visualizing higher-dimensional manifolds is a problem, and no images of three-dimensional (un)stable manifolds have been published.

The conceptual idea behind the algorithm presented here is to use the parametriza-

tion that is ideal from the geometrical point of view, irrespective of the dynamics given by the flow on $W^u(x_0)$. By this we mean the one-parameter parametrization whose elements are ‘the best possible topological $(k-1)$ -spheres’. These are given by the spheres that consist of points with constant geodesic distance d_g inside $W^u(x_0)$ from x_0 . The geodesic distance $d_g(x, y)$ is the arclength of the shortest path in $W^u(x_0)$, called a *geodesic*, connecting x and y , that is,

$$d_g(x, y) := \min_{\gamma} \left\{ \int_{\gamma} ds \mid \gamma \in W^u(x_0) \text{ is a path connecting } x \text{ with } y \right\}. \quad (2.2)$$

This gives rise to the geodesic parametrization

$$\{S_{\eta}\}_{\eta>0} \text{ where } S_{\eta} := \{x \in W^u(x_0) \mid d_g(x, x_0) = \eta\} \quad (2.3)$$

mentioned before. This parametrization does not depend on the dynamics on $W^u(x_0)$, but only on its geometry.

Because $W^u(x_0)$ is a smooth manifold tangent to $E^u(x_0)$ at x_0 , there exists a $0 < \eta_{\max} \leq \infty$ such that S_{η} is a single smooth topological $(k-1)$ -sphere without self-intersections for all $0 < \eta < \eta_{\max}$ [32]. Our method (see Section 3) can compute $W^u(x_0)$ up to $\eta_{\max}$. However, it is possible that $\eta_{\max}$ is finite. A trivial case is that the manifold converges to a regular attractor, such as an equilibrium or periodic orbit. Our implementation deals with this case by growing the manifold with different speed in different directions, as is explained in [22].

More interesting is the situation that $\eta_{\max} < \infty$ even when the manifold does not converge to an attractor. In the case $k = 2$ this occurs when for some η_c the circle S_{η_c} self-intersects and then splits up into two disjoint smooth closed curves. This means that there is a point $y \in W^u(x_0)$ with two different geodesics along which the geodesic distance η_c is realized. We did not know any example of this situation, even though some manifolds are very complicated surfaces, such as the stable manifold of the origin in the Lorenz system; see Section 4.5.

We constructed an example of a manifold with $\eta_{\max} < \infty$ and indeed our algorithm stops when $\eta_{\max}$ is reached; see Section 4.4 and Figure 4.6. However, we repeat that this geometric obstruction does not seem to occur in known applications. Moreover, it should be possible to restart the computation from the two disjoint circles; see again Section 4.4.

In fact, all methods that try to grow the manifold from an initial $(k-1)$ -sphere outward in (approximately) constant steps, including those in [10, 17, 34], encounter this geometric obstruction. It is the price for abandoning the parametrization $\{\phi^t(\mathcal{S}_{\delta})\}_{t \in \mathbb{R}}$, whose every leaf is a smooth topological $(k-1)$ -sphere in $W^u(x_0)$, in favor of the geodesic parametrization $\{S_{\eta}\}_{\eta>0}$, whose leaves naturally induce a regular discretization of $W^u(x_0)$ with a well-controlled error. We make use of this property of the geodesic parametrization in Section 5.

3. The general algorithm. The setup of the algorithm in [22] for the special case $k = 2$ and $n = 3$ generalizes to arbitrary k and n . In particular, as we will explain now, it is still possible to find a new mesh point by continuation of a one-parameter family of well-posed boundary value problems.

As before, let M_i denote the set of mesh points at step i . From the mesh points of M_i we form the simplicial complex C_i , which consists of $(k-1)$ -simplices (lines for a 1-sphere, triangles for a 2-sphere, etc.; see [29]). Hence, C_i is a piecewise linear continuous approximation of the $(k-1)$ -sphere S_{η_i} at step i . A first piece of the

manifold $W^u(x_0)$ (up to a prescribed geodesic distance $D = \eta_l$) is approximated by the k -dimensional simplicial complex $\mathcal{C}$ that is formed from the total mesh $\mathcal{M} = \cup_{0 \leq i \leq l} M_i$ in such a way that the $(k-1)$ -simplices in C_i appear as faces of the k -simplices of $\mathcal{C}$.

As starting data M_0 and C_0 we choose a mesh representation of a $(k-1)$ -sphere S_δ in $E^u(x_0)$ at some prescribed distance δ from x_0 . The algorithm now proceeds by adding new $(k-1)$ -spheres in steps. Suppose that we have computed M_i and C_i , so that we are at step k where we want to find M_{i+1} and C_{i+1} .

3.1. Finding a new point in M_{i+1} . For every point $r \in M_i$ we define an $(n-k+1)$ -dimensional hyperplane $\mathcal{F}_r$ with the property that it is ‘most perpendicular’ to C_i at r . By this we mean that we define $\mathcal{F}_r$ to be perpendicular to a suitable average of the vectors between r and its direct neighbors in M_i . Independent of the dimension k of $W^u(x_0)$, the intersection of $W^u(x_0)$ with $\mathcal{F}_r$ in the n -dimensional state space is generically a one-dimensional curve; certainly at least locally near r this curve is well-defined.

Consider now the following one-parameter family, parametrized by the integration time τ , of two-point boundary value problems finding the orbits $\{ \phi^t(q_r(\tau)) \mid t \in [0, \tau] \}$ satisfying

$$\phi^0(q_r(\tau)) = q_r(\tau) \in C_i, \quad (3.1)$$

$$\phi^\tau(q_r(\tau)) = b_r(\tau) \in \mathcal{F}_r. \quad (3.2)$$

We denote the initial condition in C_i by $q_r(\tau)$, and the final point in $\mathcal{F}_r$ by $b_r(\tau)$. Then the intersection curve $W^u(x_0) \cap \mathcal{F}_r$ can be represented in good approximation (and locally near r) as the one-parameter family $b_r(\tau)$, where $b_r(0) = r$.

What we need to find is the point $b_r = b_r(\tau_r)$ defined uniquely by the property that τ_r is the smallest integration time τ for which $\|b_r(\tau) - r\| = \Delta$. To find b_r we start from the trivial solution $q_r(0) = b_r(0) = r$ for $\tau = 0$ which clearly satisfies the boundary conditions (3.1) and (3.2). We then continue this solution in the continuation parameter τ while monitoring the test function

$$\Delta - \|b_r(\tau) - r\|. \quad (3.3)$$

When we find the first zero of (3.3) for some $\tau = \tau_r$ the continuation stops and we set $b_r = b_r(\tau_r)$.

This idea is currently only implemented for $k = 2$ using a shooting approach by searching along C_i , which is a piecewise linear curve. For $k = 3$ this would lead to a search for $q_r(\tau)$ on a set of triangles; such a two-dimensional search has already been implemented in [20] in the setting of discrete dynamical systems. A genuine boundary value problem continuation will have to deal with the fact that the initial condition (3.1) is only continuous at mesh points; this is work in progress. Note that for $k > 3$ the boundary value problem continuation is still well-defined as a one-parameter continuation; $q_r(\tau)$ would be traced ‘automatically’ along a curve in C_i by the boundary value solver.

3.2. Checking the guess for Δ_i . In order to decide whether Δ_i was appropriate we consider for each $r \in M_i$ the curvature of the one-dimensional curve through r and the corresponding points $b_r \in M_{i+1}$ and $p_r \in M_{i-1}$. This is motivated by what is done for one-dimensional manifolds in [21] and goes back to an idea of Hobson [14]. Let α_r denote the angle between the line through p_r and r and the line through r and

b_r . The step Δ_i was acceptable if both

$$\alpha_r < \alpha_{\max}, \quad (3.4)$$

$$\Delta_i \cdot \alpha_r < (\Delta\alpha)_{\max} \quad (3.5)$$

hold for all $r \in M_i$. In this case we accept M_{i+1} and form the simplicial complex C_{i+1} as the next $(k-1)$ -sphere. On the other hand, if there is some $r \in M_i$ such that either (3.4) or (3.5) is not satisfied then Δ_i was too big. We discard M_{i+1} , decrease Δ_i (in practice we halve it), and compute a new M_{i+1} at this smaller distance from M_i . The algorithm accepts M_{i+1} after all (with a warning message), when Δ_i has already been decreased to a prespecified minimal value of $\Delta_{\min}$.

If Δ_i was acceptable, we usually set $\Delta_{i+1} = \Delta_i$. However, if for every $r \in M_i$ both α_r and $\Delta \cdot \alpha_r$ are well below the respective upper bounds in (3.4) or (3.5), say, less than $\alpha_{\min}$ and $(\Delta\alpha)_{\min}$ respectively, then we try a larger Δ_{i+1} in the next step (in practice we set $\Delta_{i+1} = 2\Delta_i$). The parameters $\alpha_{\min}$, $\alpha_{\max}$, $(\Delta\alpha)_{\min}$, and $(\Delta\alpha)_{\max}$ need to be specified by the user before a computation.

3.3. Adding and removing mesh points. It is very important to maintain an accurate approximation of S_{i+1} by the simplicial complex C_{i+1} , because this allows us to control the interpolation error; see Section 5. With our method for controlling the growth step size Δ_i we ensure that the distance between C_i and C_{i+1} is within the accuracy bounds of the computation. However, we also need to make sure that, as neighboring points of M_{i+1} are too far away from each other inside this level set, a new point is added between them. To do this we chose a suitable point in C_i between the two respective points in M_i and then use the method in Section 3.1 to find a new point in C_{i+1} . In this way, we ensure that no interpolation is performed between points that are further away from each other than a maximally allowed and prespecified distance $\Delta_{\mathcal{F}}$.

Similarly, when two points in M_{i+1} become too close, we remove one of them to ensure proper order relations between directly neighboring points.

3.4. Implementation details for $k = 2$. While the general setup remains the same, there are serious difficulties of implementing the general algorithm for $k \geq 3$. Already for $k = 3$, it is a nontrivial task to construct the simplicial complexes C_i from the sets of mesh points M_i , and then assemble them to obtain the approximation $\mathcal{C}$; see, for example, [29]. Furthermore, developing a general strategy for adding and removing mesh points to ensure a uniform mesh $\mathcal{M}$ is very difficult in this case.

These are the main reasons why our general algorithm, or any other growth algorithm for that matter, has only been implemented for two-dimensional manifolds. The algorithm is presently implemented for manifolds of dimension $k = 2$ in a state space of arbitrary dimension n . This includes the case of an unstable manifold of a periodic orbit of saddle type, for which the starting circle S_δ needs to be chosen in the linear unstable eigenspace of the periodic orbit; see Section 4.3 and [25, 26] for details.

For $k = 2$ the manifold $W^u(x_0)$ is parametrized by a family of (topological) circles S_η . The 1-simplices that form the continuous objects C_i are line segments between neighboring points in M_i . The algorithm grows $W^u(x_0)$ by adding a new circle C_{i+1} at each step, which means that a new shell or band of width Δ_i of triangles is added to $\mathcal{C}$. The simplicial complex $\mathcal{C}$ representing $W^u(x_0)$ consists of triangles that constitute the shells between consecutive circles C_{i-1} and C_i .

FIG. 4.1. The stable manifold $W^s(0)$ of the Lorenz system for $\sigma = 10$, $\varrho = 28$, and $\beta = 0.4$, computed up to geodesic distance 66.28 (a). Close-ups near the attracting periodic orbit show the growth process: 24 bands or geodesic distance 40.25 (b), 34 bands or geodesic distance 46.25 (c), and 49 bands or geodesic distance 54.75 (d).

The initial circle M_0 is chosen in $E^u(x_0)$ at distance δ . As was explained in Section 3.2, the width of each ring depends on the curvature of the manifold locally near the ring, and it is governed by $\alpha_{\min}$, $\alpha_{\max}$, $(\Delta\alpha)_{\min}$, and $(\Delta\alpha)_{\max}$. Extra points are added to M_i when neighboring mesh points are further apart than $\Delta_{\mathcal{F}}$, and removed if they are closer than $\delta_{\mathcal{F}}$. The implementation of this is quite straightforward for $k = 2$ because of the order structure of the meshes M_i .

The boundary value problem (3.1)-(3.2) is presently solved with a shooting approach, for which we use a fourth-order fixed time-step Runge-Kutta integration routine.

4. Examples. We now present a number of examples to discuss the properties of our algorithm. In particular, we explain how the parametrization of the manifolds by arclength can be used to gain insight into their often quite complicated geometry. All figures have been rendered with the program Geomview [28], and each figure is accompanied by an animation that is supplied with this paper.

4.1. The Lorenz system with an attracting orbit. We start with an example of a complicated two-dimensional stable manifold of the origin in the well-known Lorenz system [19], but for a non-standard value of the parameter β . The Lorenz system is defined as

$$\begin{cases} \dot{x} &= \sigma(y - x), \\ \dot{y} &= \varrho x - y - xz, \\ \dot{z} &= xy - \beta z. \end{cases} \quad (4.1)$$

For the choice of parameters $\sigma = 10$, $\varrho = 28$, and $\beta = 0.4$ there is an attracting periodic orbit that encircles two saddle points. Figure 4.1 (a) shows the two-dimensional stable manifold $W^s(0)$ of the origin.

The manifold $W^s(0)$ was computed with the following accuracy. The computation started with 20 points on a circle in $E^s(0)$ of radius $\delta = 1.0$ around the origin. Then new circles were added at distances Δ_i controlled by $\alpha_{\min} = 0.3$, $\alpha_{\max} = 0.4$, $(\Delta\alpha)_{\min} = 0.1$, and $(\Delta\alpha)_{\max} = 1.0$; see Sections 3.1 and 3.2. (For practical reasons a new circle was always accepted if $\Delta \leq 0.01$.) The mesh points on a circle are never more than $\Delta_{\mathcal{F}} = 1.0$ and less than $\delta_{\mathcal{F}} = 0.25$ apart. In total 75 circles were computed on $W^s(0)$; the last circle is approximately at geodesic distance 66.28 from the origin, and it consists of 1522 mesh points. How $W^s(0)$ is grown during the computation is shown in Figure 4.1 (b)–(d) with three close-up views near the attracting periodic orbit; see also the accompanying movie. We remark that the manifold in Figure 4.1 (a) is featured on the advertising poster for SIADS.

The dark and light bands indicate the positions of the rings that are added during the course of the computation. They illustrate the geodesic level sets on $W^u(0)$ that were computed. To highlight this, Figure 4.2 shows individual small segments of bands at different (approximately) constant geodesic distance from the origin. Each boundary consists of two rings that are both indeed smooth, simple and closed. Notice that the rings have quite complicated embeddings in the phase space due to the overall shape of the manifold, but are nevertheless topologically trivial, that is, contractable to the fixed point from which they were grown. For an animation illustrating the change in shape of the bands and rings during the growth process see the accompanying movie.

4.2. An optimal control system. The following example from optimal control theory is taken from [16] and demonstrates that our implementation for $k = 2$ can indeed be used in ambient spaces of arbitrary dimension. Furthermore, it is an illustration of how invariant manifold computations can be used to understand the dynamics of systems arising in applications.

Consider an inverted planar pendulum balancing on a cart. The cart moves in the plane of the pendulum with an applied horizontal force u constituting a control. The mass of the cart is M , the mass of the pendulum m and its center of mass is at distance l from the pivot. Disregarding the model equations associated with the cart

FIG. 4.2. Growing bands are approximations of the level sets S_η of the stable manifold of the Lorenz system in Figure 4.1. From (a) to (d) are shown band 14 with $\eta \in [30.25, 32.25]$, bands 23–25 with $\eta \in [38.25, 41.25]$, bands 33–35 with $\eta \in [44.75, 47.25]$, and bands 68–71 with $\eta \in [63.75, 65.75]$.

gives the two-dimensional vector field

$$\begin{cases} \dot{x}_1 &= x_2, \\ \dot{x}_2 &= f(x_1, x_2) + c(x_1, x_2)u \\ &:= \frac{\frac{g}{l} \sin(x_1) - \frac{1}{2} m_r x_2^2 \sin(2x_1) - \frac{m_r}{ml} \cos(x_1)u}{\frac{4}{3} - m_r \cos^2(x_1)}. \end{cases} \quad (4.2)$$

Here x_1 is the angle measured from the vertical up position, $m_r = m/(m + M)$ is the mass ratio, and g is the gravitational constant. Note that the origin is an unstable equilibrium corresponding to the vertical position above the pivot.

We wish to find an optimal control u as a function of x_1 and x_2 that drives the

FIG. 4.3. Four views of the two-dimensional stable manifold $W^s(0)$ up to geodesic distance 16.5 of the Hamiltonian system (4.4) projected onto three-dimensional spaces, namely onto $\{p_2 = 0\}$ (a), onto $\{p_1 = 0\}$ (b), onto $\{x_2 = 0\}$ (c), and onto $\{x_1 = 0\}$ (d). The self-intersections in (d) are due to the projection.

system to the origin, while minimizing the cost function

$$Q(x_1, x_2, u) = \mu_1 x_1^2 + \mu_2 x_2 + \mu_u u^2, \quad (4.3)$$

for positive parameters μ_1 , μ_2 and μ_3 . Pontryagin's maximum principle [33] ensures that an optimal solution exists and is represented by special solutions of the four-

dimensional Hamiltonian system given by the Hamiltonian:

$$H(x_1, x_2, p_1, p_2) = Q(x_1, x_2, u^*(x_1, x_2, p_1, p_2)) + p_1 x_2 + p_2 f(x_1, x_2) + p_2 c(x_1, x_2) u^*(x_1, x_2, p_1, p_2), \quad (4.4)$$

where $u^*(x_1, x_2, p_1, p_2) = -\frac{1}{2\mu_u} c(x_1, x_2) p_2$. One is particularly interested in the geometrical properties of the two-dimensional stable manifold of the origin, because $u^*(x_1, x_2, p_1, p_2)$ is the optimal control achieving the minimum for the cost function (4.3) provided (x_1, x_2, p_1, p_2) lies on $W^s(0)$. If there is more than one point on $W^s(0)$ for given x_1 and x_2 , then there exists more than one choice for $u^*(x_1, x_2, p_1, p_2)$. This means that suboptimal (only locally optimal) solutions exist or that the globally optimal cost can be achieved using two different control actions.

We computed $W^s(0)$ using the parameter values as in [16], namely $m = 2$ kg, $M = 8$ kg, $l = 0.5$ m, $g = 9.8$ m/s², $\mu_1 = 0.1$, $\mu_2 = 0.05$ and $\mu_u = 0.01$. We used the following accuracy parameters for the computation. The starting data was a set of 20 points on a circle in $E^s(0)$ of radius $\delta = 1.0$ around the origin. The distance Δ at which new circles were added was controlled by $\alpha_{\min} = 0.3$, $\alpha_{\max} = 0.4$, $(\Delta\alpha)_{\min} = 0.05$, $(\Delta\alpha)_{\max} = 0.2$, where we always accepted the circle if $\Delta \leq 0.05$; see Sections 3.1 and 3.2. The distances between mesh points on a circle were always in between $\delta_{\mathcal{F}} = 0.125$ and $\Delta_{\mathcal{F}} = 0.5$.

The result is shown in Figure 4.3 where $W^s(0)$ was computed up to a geodesic distance of approximately 16.5, corresponding to 65 circles. The last circle has 481 mesh points. Because the ambient space is four-dimensional it is difficult to view the result. Therefore, Figure 4.3 shows four views in four different three-dimensional projections. The growth process of the manifold is shown in the same projections in the accompanying animation. Note that the optimal control $u^*(x_1, x_2, p_1, p_2) = -\frac{1}{2\mu_u} c(x_1, x_2) p_2$ does not depend on p_1 . Therefore, the folds shown in the projection onto $\{p_1 = 0\}$ in Figure 4.3 (b) indicate that indeed more than one optimal control exists for a set of (x_1, x_2) values.

4.3. A Möbius strip in the ζ^3 -model. To illustrate that the algorithm is able to compute stable or unstable manifolds of periodic orbits of saddle-type we consider the vector field

$$\begin{cases} \dot{x} &= y, \\ \dot{y} &= z, \\ \dot{z} &= \alpha x - x^2 - \beta y - z. \end{cases} \quad (4.5)$$

It is known as the ζ^3 -model [2], and arises as the asymptotic normal form of a system near the simultaneous onset of up to three instabilities; see also [17, 18, 22]. For $\alpha = 3.2$ and $\beta = 2$ the system has a saddle periodic orbit Γ with Floquet multipliers 1, -0.022 , and -1.15 , approximately. This means that both the stable and unstable manifolds are non-orientable.

We computed $W^s(\Gamma)$, where the accuracy parameters were $\Delta_{\mathcal{F}} = 0.1$, $\delta_{\mathcal{F}} = 0.025$, $\Delta_{\min} = 0.01$, $0.2 < \alpha < 0.3$, and $0.001 < \Delta\alpha < 0.1$. We took 100 mesh points on Γ and the initial circle in the linear unstable eigenspace $N^s(\Gamma)$ of Γ was a double cover with 200 points at distance $\delta = 0.1$ from Γ . More specific details on how to compute non-orientable manifolds can be found in [25]. The result is shown in Figure 4.4 for a total computed geodesic distance of 3.1, corresponding to 100 circles. To help with depth perception we colored the last ring differently. It forms the boundary of the computed piece of the manifold, which is indeed a Möbius strip; see also the accompanying animation of the growth process.

FIG. 4.4. Two views of the two-dimensional stable manifold up to geodesic distance 3.1 of the periodic orbit of the ζ^3 model (4.5) for $\alpha = 3.2$ and $\beta = 2$. The manifold is in fact a Möbius strip; to see this its edge is highlighted.

A saddle periodic orbit in $\mathbb{R}^3$ that lost its stability in a period-doubling has a nonorientable unstable manifold $W^u(\Gamma)$ whose boundary is the nearby period-doubled attracting periodic orbit. Its stable manifold $W^s(\Gamma)$ is necessarily also nonorientable, but it does not need to have a boundary. The stable nonorientable manifold in Figure 4.4 extends all the way to infinity, folding back and forth through the ‘hole’ bounded by the periodic orbit Γ . Nevertheless, $W^s(\Gamma)$ does not separate the phase space into two invariant subsets. More details of how nonorientable manifolds permit insight into dynamical systems can be found in the survey paper [26].

4.4. A manifold with $\eta_{\max} < \infty$. As was mentioned in Section 2, an unstable manifold $W^u(x_0) = \{S_\eta\}_{\eta>0}$ may be such that $\eta_{\max} < \infty$, even though the manifold does not converge to an attractor. This means that $S_{\max}$ is the first element of the parametrization that ceases to be a smooth topological $(k-1)$ -sphere without self-intersections. We are not aware of any example of this phenomenon in a dynamical system, and so we constructed an example of this situation for $k = 2$ as follows. Consider the vector field

$$\begin{cases} \begin{pmatrix} \dot{u} \\ \dot{v} \end{pmatrix} = A \begin{pmatrix} u \\ v \end{pmatrix} \\ \dot{w} = -w + u^2 + v^2, \end{cases} \quad (4.6)$$

where

$$A = \begin{pmatrix} \alpha & \omega \\ -\omega & \alpha \end{pmatrix}.$$

For $\alpha > 0$ the origin of this vector field has a parabola-shaped unstable manifold tangent to the (x, y) -plane. The dynamics on this unstable manifold is simply given by the constant rotation $\omega > 0$.

FIG. 4.5. The computed unstable manifold of the origin of the constructed vector field of Section 4.4 with a typical orbit (in green) on it for $a = 5$ (a). Panel (b) shows the mesh resulting from the computation.

We now apply the coordinate transformation

$$\begin{cases} x &= u - w \\ y &= v/3 \\ z &= (w + a u^2) \end{cases} \quad (4.7)$$

to (4.6), where $a > 0$ is a parameter. This deforms the parabola-shaped unstable manifold of (4.6) first in the x -direction so that only a finite piece of it lies to the right of the (y, z) -plane. Furthermore, for sufficiently large a a hump is created on the manifold in the region of positive x . The shape of the manifold is independent of ω , the amount of spiralling on it. We took $\omega = 0.2$ in the computations of $W^u(x_0)$.

Figure 4.5 shows the resulting manifolds for fixed $\alpha = 0.1$ and $a = 5$. For sufficiently large ω one can get a good impression of the shape of the manifold simply by plotting a single orbit on $W^u(0)$; this is illustrated in Figure 4.5 (a) for an initial condition $(0.005, 0, 0)$ in $E^u(x_0)$ with $\omega = 10$. Our algorithm is able to compute $W^u(0)$ for $a = 5$ but it makes quite small Δ_i steps near and just after the ‘bump’. Notice that quite a number of mesh points need to be removed before the mesh is restored such that the manifold can be grown further with larger steps, as is further illustrated in the accompanying animation.

In Figure 4.6 for $a = 10$ on the other hand, the ‘bump’ is now much larger. The trajectory in Figure 4.6 (a) was again computed with $\omega = 10$ and the same initial condition as for $a = 5$. There exists a point at the bottom of the ‘valley’ between the infinite part and the bump on the manifold that has two geodesics. In other words, it is a singularity of the geodesic level set on which it lies, which forms a figure eight and then splits into two disjoint circles, one around the bump and one around the ‘neck’ of the manifold. In other words, the valley point is indeed a point whose geodesic distance is $\eta_{\max}$ as defined above. Indeed the algorithm stops when the valley floor is reached; see also the accompanying animation.

FIG. 4.6. The computed unstable manifold of the origin of the constructed vector field of Section 4.4 with a typical orbit (in green) on it for $a = 10$ (a). Panel (b) shows the mesh that resulted from the computation. The algorithm stops when the geodesic level set self-intersects and then is no longer connected.

In both computations the starting data was a circle in $E^u(0)$ of radius $\delta = 0.01$ around the origin. Furthermore we set $\alpha_{\min} = 0.3$, $\alpha_{\max} = 0.4$, $(\Delta\alpha)_{\min} = 0.002$, $(\Delta\alpha)_{\max} = 0.015$ and $\delta_{\mathcal{F}} = 0.005$ and $\Delta_{\mathcal{F}} = 0.04$.

We remark that this problem occurs for algorithms that use the approach of calculating approximate geodesic level sets, that is, also for the algorithms in [10, 17, 34]. Actually, this geometric obstruction to growing $W^u(0)$ is, in a sense, an artificial problem. It can be overcome by changing to a different metric; in this specific example one can rescale y to ensure that the geodesic distance in the valley is realized by the path exactly over the top. However, changing the metric as one encounters this obstruction during a computation is not practical. Instead one could split up the figure-eight level set into two separate circles and start two individual computations. Since we did not encounter this geometric obstruction in any other example, this has not been implemented.

4.5. The Lorenz system with a strange attractor. To demonstrate how one can make use of the parametrization by geodesic distance to understand complicated manifolds, we consider again the Lorenz system (4.1), but now for the standard values of the parameters $\sigma = 10$, $\rho = 28$, and $\beta = 8/3$ when there is the well-known Lorenz attractor. This example has emerged as somewhat of a test case for the computation of two-dimensional global manifolds, and the task is to compute the stable manifold of the origin that spirals into the chaotic attractor. How exactly does this happen and what does the manifold look like?

To begin with, any piece of $W^s(0)$ of finite geodesic distance D from the origin is simply the image of a two-dimensional disk under a smooth map. This map is implicitly given by mapping the circles and radii of the disk to the geodesic level sets and geodesics on $W^s(0)$, respectively. It is truly astounding how complicated this embedding and the image $W^s(0)$ are. A first attempt at visualizing this manifold was

FIG. 4.7. The stable manifold $W^s(0)$ of the Lorenz system for $\sigma = 10$, $\varrho = 28$, and $\beta = 8/3$, computed up to geodesic distance 151.75 (a). Close-ups near the z -axis show the growth process: 25 bands or geodesic distance 62.75 (b), 40 bands or geodesic distance 92.75 (c), and 55 bands or geodesic distance 122.75 (d). Color indicates geodesic distance from the origin as given by the color bar.

a series of three sketches in [1]. Computed images of differing sophistication can be found in [9, 10, 13, 20, 22] and in [23], which features a direct comparison with the sketches in [1].

The manifold $W^s(0)$ in Figure 4.7 (a) was computed with the following accuracy. The computation started with 20 points on a circle in $E^s(0)$ of radius $\delta = 1.0$ around the origin. Then new circles were added at distances Δ controlled by $\alpha_{\min} = 0.3$, $\alpha_{\max} = 0.4$, $(\Delta\alpha)_{\min} = 0.1$, and $(\Delta\alpha)_{\max} = 1.0$; see Sections 3.1 and 3.2. (For practical reasons a new circle was always accepted if $\Delta \leq 0.01$.) The mesh points on a circle are never more than $\Delta_{\mathcal{F}} = 1.0$ and less than $\delta_{\mathcal{F}} = 0.25$ apart. In total 72 circles were computed on $W^s(0)$; the last circle is approximately at geodesic distance 151.75 from the origin, and it consists of 2300 mesh points.

FIG. 4.8. Growing bands are approximations of the level sets S_η of the stable manifold of the Lorenz system in Figure 4.7. From (a) to (f) are shown bands 24–25 with $\eta \in [58.75, 62.75]$, bands 30–31 with $\eta \in [70.75, 74.75]$, bands 52–54 with $\eta \in [114.75, 120.75]$, bands 56–58 with $\eta \in [122.75, 128.75]$, bands 60–62 with $\eta \in [130.75, 136.75]$, and bands 67–71 with $\eta \in [144.75, 151.75]$.

FIG. 4.9. The stable manifold $W^s(0)$ in Figure 4.7 (a) where only every second band is shown.

How $W^s(0)$ is grown during the computation is shown in Figure 4.7 (b)–(d) with three close-ups showing how a helix develops near the positive z -axis; see also the accompanying movie. The coloring indicates geodesic distance from the origin, as given by the color bar. Also shown is the one-dimensional unstable manifold of the origin in red, whose closure forms the well-known Lorenz attractor, an orbit on which is shown in yellow.

The geometry of the manifold is illustrated in Figure 4.8 and the accompanying movie by following the development of a small set of bands as they grow. This gives an idea of how $W^s(0)$ is parametrized by geodesic level sets. In spite of the complicated geometry, each of the level sets is still a nice smooth topological circle, albeit not a very ‘circular’ and flat one. The later bands nicely show how the manifold has to cross over and back between the two scrolls of the chaotic attractor. Clearly visible in Figure 4.8 (e) and (f) is the development of two further helices of $W^s(0)$ near the z -axis. Figure 4.9 shows the entire manifold rendered by plotting only every second band, to produce a see-through effect. The manifold is now plotted without the chaotic attractor and spins around in the accompanying movie.

Finally, in Figure 4.10 we show the part of the stable manifold $W^s(0)$ (together with the unstable manifold and the Lorenz attractor) inside a sphere around the point

FIG. 4.10. The stable manifold $W^s(0)$ in Figure 4.7 (a) inside a sphere of radius 30 (a), 55 (b), and 65 (c) around the point $(0,0,27)$.

(0,0,27) of radius 30, 55 and 65 in panels (a), (b) and (c), respectively. (The accompanying movie shows an animation where this radius varies.) The images on the right again show only every second computed band. This visualization aims to highlight how the manifold spirals into the chaotic attractor. Here, the coloring with geodesic distance is particularly helpful: points with the same color lie at the same geodesic distance from the origin, where red is furthest. One of the extra helixes near the z -axis is clearly visible in Figure 4.10 (b).

5. Convergence of the algorithm. Suppose that we computed a simplicial complex $\mathcal{C}(\delta, \Delta)$ with the mesh points $\mathcal{M}(D, \delta, \Delta) = \cup_{i=0}^l M_i$ approximating the first piece $\mathcal{S}_D = \{S_\eta\}_{0 \leq \eta \leq D}$ of $W^u(x_0)$ of fixed geodesic distance D . In the discussion of the convergence of the algorithm, we assume that $D \ll \eta_{\max}$, for example, $\eta_{\max} = \infty$. The dependence on δ and Δ indicates that the approximation was computed with these prespecified accuracy parameters, where for simplicity we chose $\Delta_{\mathcal{F}} = \Delta$ and $\Delta_i = \Delta$ constant so that Δ specifies the overall mesh quality.

We want to show that $\mathcal{C}(\delta, \Delta)$ converges to $\mathcal{S}_D$ as the initial step δ from x_0 and the mesh quality parameters Δ go to zero. The fundamental difficulty with any such proof is that when δ and Δ are changed, the entire approximation needs to be recomputed, meaning that the refined set of mesh points $\mathcal{M}(\tilde{\delta}, \tilde{\Delta})$ generally does not have any mesh points in common with $\mathcal{M}(\delta, \Delta)$. This is why one cannot use strategies of proof as they are used when proving the correctness of computations of implicitly defined manifolds (given as the zero set of a function) where it is possible to refine and improve the mesh locally; see, for example, [29]. In fact, the problem here is quite similar to proving the convergence of numerical approximations of attractors to the actual attractor of the underlying continuous dynamical system; see [31].

The key problem is to control the total (interpolation) error between mesh points uniformly over the mesh. We use here the approach in [20] for two-dimensional unstable manifolds of maps.

Our goal is to prove that for any given $\varepsilon > 0$ the approximation $\mathcal{C}(\delta, \Delta)$ is in an ε -neighborhood of $W^u(x_0)$, provided that δ and Δ are chosen small enough. This is known as upper semicontinuity. Our proof essentially tracks how different errors grow along the manifold, which is described by certain Lipschitz constants. The key problem is to control the total (interpolation) error between mesh points uniformly over the mesh.

It is a fact of life that the Lipschitz constants generally cannot be determined in a given system. Therefore, it is impossible to derive *a priori* error bounds for the computation of global invariant manifolds. The only practical way of checking a computation is to repeat it with increased accuracy and compare the results. Generally speaking, the longer the piece of the manifold one wants to compute, that is the larger D , the smaller one should choose the accuracy parameters.

5.1. Upper semicontinuity. Recall (see, for example, [15]) that the semi-distance between two sets $A, B \subset \mathbb{R}^n$ is

$$d(A, B) = \sup_{y \in A} \{d(x, B)\}, \text{ where}$$

$$d(x, B) = \inf_{y \in B} \{\|x - y\|\}.$$

Let $\mathcal{N}_\varepsilon(X)$ denote the ε -neighborhood of a set $X \subset \mathbb{R}^n$.

THEOREM 1 (Upper semicontinuity).

Let $D > 0$ be given and fixed, and let $\mathcal{C}(\delta, \Delta)$ denote the simplicial complex approxi-

ation of $\mathcal{S}_D = \{S_\eta\}_{0 \leq \eta \leq D}$. Then for any $\varepsilon > 0$

$$\mathcal{C}(\delta, \Delta) \subset \mathcal{N}_\varepsilon(W^u(x_0))$$

for all sufficiently small δ and Δ .

Proof:

In order to show that $\mathcal{C}(\delta, \Delta) \subset \mathcal{N}_\varepsilon(W^u(x_0))$ we will show that $\mathcal{C}(\delta, \Delta) \subset \mathcal{N}_\varepsilon(\mathcal{S}_{\tilde{D}})$ for some $\tilde{D} > D$. Since $\mathcal{S}_{\tilde{D}} \subset W^u(x_0)$ this will automatically imply that $\mathcal{C}(\delta, \Delta) \subset \mathcal{N}_\varepsilon(W^u(x_0))$. The reason for using $\mathcal{S}_{\tilde{D}}$ instead of $W^u(x_0)$ is that $\mathcal{S}_{\tilde{D}}$ is compact, which will be used to get Lipschitz bounds.

The total error $\varepsilon_{\mathcal{C}}$ is defined as

$$\varepsilon_{\mathcal{C}} = d(\mathcal{C}(\delta, \Delta), \mathcal{S}_{\tilde{D}}). \quad (5.1)$$

This means that $\mathcal{C}(\delta, \Delta)$ is contained in $\mathcal{N}_\varepsilon(\mathcal{S}_{\tilde{D}})$ for any $\varepsilon > \varepsilon_{\mathcal{C}}$. We obtain upper semicontinuity by showing that $\varepsilon_{\mathcal{C}} \rightarrow 0$ as $\delta, \Delta \rightarrow 0$.

There are two contributions to the error $\varepsilon_{\mathcal{C}}$. First, the mesh points in $\mathcal{M}(\delta, \Delta)$ do not lie exactly on $\mathcal{S}_{\tilde{D}}$, which gives rise to the mesh error

$$\varepsilon_{\mathcal{M}} = d(\mathcal{M}(\delta, \Delta), \mathcal{S}_{\tilde{D}}).$$

Second, there is the global interpolation error $\varepsilon_I(\Delta)$ due to the fact that points in $\mathcal{C}(\delta, \Delta) \setminus \mathcal{M}(\delta, \Delta)$ are interpolated linearly between the mesh points. Consequently, we have for the total error

$$\varepsilon_{\mathcal{C}} \leq \varepsilon_{\mathcal{M}} + \varepsilon_I(\Delta). \quad (5.2)$$

Note that, unlike in the error analysis of the approximation of an implicitly defined invariant manifold [29], these two errors are not independent. The interpolation error at step i enters into the mesh error at step $i + 1$. Hence, if we wish to compute more accurate mesh points we need to start an entirely new computation with higher accuracy.

Clearly the interpolation error $\varepsilon_I(\Delta)$ goes to zero with Δ , so that in light of (5.2) it suffices to show that $\varepsilon_{\mathcal{M}}$ goes to zero with Δ and δ . To this end, we define the mesh error at step i

$$\varepsilon_{\mathcal{M}}(i) = d(M(\delta, \Delta)_i, \mathcal{S}_{\tilde{D}}).$$

The initial mesh error $\varepsilon_{\mathcal{M}}(0)$ is due to the fact that $M(\delta, \Delta)_0$ is chosen in the linear subspace $E^u(x_0)$ at distance δ from x_0 , instead of on $W^u(x_0)$. By taking δ small, $\varepsilon_{\mathcal{M}}(0)$ can be made arbitrarily small. Under the assumption that $W^u(x_0)$ is C^2 , the initial error $\varepsilon_{\mathcal{M}}(0)$ is of order $O(\delta^2)$.

The mesh error $\varepsilon_{\mathcal{M}}(i + 1)$ at step $i + 1$ can be estimated in terms of $\varepsilon_{\mathcal{M}}(i)$ as follows. Points in $M(\delta, \Delta)_{i+1}$ are found by continuing the boundary value problem (3.1)-(3.2). Since the accuracy of the boundary value solver is independent of the other computational errors, we assume that solving the boundary value problem does not introduce an extra error. (In practice, it is sufficient to ensure that the boundary value solver has an error of less than the order of the interpolation error.) Recall that, for any $r \in M(\delta, \Delta)_i$, the starting condition (3.1) of the boundary value problem leads to points $q_r(\tau_r) \in S(\delta, \Delta)_i$ and the end condition (3.2) to points $b_r(\tau_r) \in \mathcal{F}_r$, where τ_r is the integration time such that $\|b_r(\tau_r) - r\| = \Delta$. The error at $S(\delta, \Delta)_{i+1}$ comes

from the fact that $S(\delta, \Delta)_i$ is only approximately on $\mathcal{S}_{\bar{D}}$, so that we are solving the ‘wrong’ boundary value problem. Therefore,

$$\begin{aligned}\varepsilon_{\mathcal{M}}(i+1) &= \max_{r \in M(\delta, \Delta)_i} d(b_r(\tau_r), \mathcal{S}_{\bar{D}}) \\ &= \max_{r \in M(\delta, \Delta)_i} d(\phi^{\tau_r}(q_r(\tau_r)), \mathcal{S}_{\bar{D}}).\end{aligned}\quad (5.3)$$

Since $\mathcal{S}_{\bar{D}}$ is compact, we may assume that the vector field satisfies a Lipschitz condition in a neighborhood U of $\mathcal{S}_{\bar{D}}$. This means that

$$d(\phi^t(p), \mathcal{S}_{\bar{D}}) \leq e^{\kappa_U t} d(p, \mathcal{S}_{\bar{D}}) \quad (5.4)$$

for some constant $\kappa_U > 0$, and such that the orbit of p up to $\phi^t(p)$ is in U ; see also [31]. Combining this with (5.3) and defining τ_i such that $e^{\kappa_U \tau_i} = \max_{r \in M(\delta, \Delta)_i} (e^{\kappa_U \tau_r})$,

we get

$$\begin{aligned}\varepsilon_{\mathcal{M}}(i+1) &\leq \max_{r \in M(\delta, \Delta)_i} e^{\kappa_U \tau_r} d(q_r(\tau_r), \mathcal{S}_{\bar{D}}) \\ &\leq e^{\kappa_U \tau_i} \varepsilon_{\mathcal{C}}(i) \leq e^{\kappa_U \tau_i} (\varepsilon_{\mathcal{M}}(i) + \varepsilon_I(\Delta)).\end{aligned}\quad (5.5)$$

The local manifold is always an attractor [27], so there exists a neighborhood $V \subset U$ of x_0 on which we find a *negative* Lipschitz constant $-\kappa_V < 0$. By choosing δ sufficiently small we can assure that $S(\delta, \Delta)_0 \subset V$. Suppose that the first, say, J spheres lie in V , that is, $S(\delta, \Delta)_i \subset V$ for $0 \leq i \leq J$. We then have

$$\varepsilon_{\mathcal{M}}(i+1) \leq e^{-\kappa_V \tau_i} (\varepsilon_{\mathcal{M}}(i) + \varepsilon_I(\Delta)) \leq \varepsilon_{\mathcal{M}}(1) \quad \text{for } 0 \leq i < J, \quad (5.6)$$

provided Δ is small enough.

We now show that the error on $U \setminus V$ can be controlled. For the $(k-1)$ -spheres in $U \setminus V$ of which there are, say, N , using (5.5) and the fact that $\kappa_U > 0$ we get

$$\begin{aligned}\varepsilon_{\mathcal{M}}(J) &\leq \dots \leq \varepsilon_{\mathcal{M}}(N+J) \\ &\leq e^{\kappa_U \bar{\tau}} \varepsilon_{\mathcal{M}}(J) + N e^{\kappa_U \bar{\tau}} \varepsilon_I(\Delta).\end{aligned}\quad (5.7)$$

Here, the number $\bar{\tau}$ is chosen such that $\sum_{i=J}^{N+J-1} \tau_i \leq \bar{\tau}$. In fact, $\bar{\tau}$ is an upper bound on the total integration time that is needed by the algorithm to cover $U \setminus V$, which depends only on the prescribed geodesic distance D (once V is fixed). In other words, $\bar{\tau}$ can be chosen such that it is independent of N . By combining (5.6) and (5.7) we have

$$\begin{aligned}\varepsilon_{\mathcal{M}} &= \max_{0 \leq i \leq N+J} \{ \varepsilon_{\mathcal{M}}(i) \} \\ &= \max \{ \varepsilon_{\mathcal{M}}(0), \varepsilon_{\mathcal{M}}(N+J) \} \\ &\leq e^{\kappa_U \bar{\tau}} \varepsilon_{\mathcal{M}}(0) + N e^{\kappa_U \bar{\tau}} \varepsilon_I(\Delta).\end{aligned}\quad (5.8)$$

The term $e^{\kappa_U \bar{\tau}} \varepsilon_{\mathcal{M}}(0)$ in (5.8) can be made arbitrarily small by decreasing δ . In order to make the term $N e^{\kappa_U \bar{\tau}} \varepsilon_I(\Delta)$ in (5.8) arbitrarily small $\varepsilon_I(\Delta)$ needs to decrease faster than linear with Δ . This is guaranteed by the assumption that $W^u(x_0)$ is at least C^2 , in which case $\varepsilon_I(\Delta) = O(\Delta^2)$. It follows that $\varepsilon_{\mathcal{C}} \rightarrow 0$ as $\delta, \Delta \rightarrow 0$, which concludes the proof. $\square$

REMARK 2.

In practice, we do vary Δ_i during a computation. Since the variation of Δ_i is related to local curvature along geodesics, the overall interpolation error does not increase; see [14, 21]. Furthermore, we add and remove mesh points, but any required interpolation is done between points at most $\Delta_{\mathcal{F}}$ apart; see Section 3.3.

5.2. Upper semicontinuity. The main difficulty in proving that $\mathcal{C}(\delta, \Delta)$ converges to $\mathcal{S}_D$ (in the Hausdorff metric) lies in showing that the last $(k - 1)$ -sphere of $\mathcal{C}(\delta, \Delta)$ actually converges to the geodesic level set $\mathcal{S}_D$. The problem is that the entire mesh changes as δ or Δ vary. Hence, it is not possible to consider the convergence of a specific mesh point and its geodesic distance. However, we believe this to be a technical difficulty and we have the following.

CONJECTURE 3 (Lower semicontinuity).

For given fixed $D > 0$ let $\mathcal{C}(\delta, \Delta)$ be the simplicial complex approximating of $\mathcal{S}_D$. Then for any $\varepsilon > 0$

$$\mathcal{S}_D \subset \mathcal{N}_\varepsilon(\mathcal{C}(\delta, \Delta))$$

for all sufficiently small δ and Δ .

If we assume that the vector field, and hence $W^u(x_0)$, is at least C^2 , then we can conclude that the algorithm does not bypass a part of $\mathcal{S}_D$. This is the case, because by Theorem 1 $\mathcal{C}(\delta, \Delta)$ lies in an ε -neighborhood of the C^2 manifold $\mathcal{S}_D$. Consequently, any ‘bump’ is C^2 , and $\mathcal{C}(\delta, \Delta)$ must stay ε -close to the bump for δ and Δ small enough.

From a practical point of view this is the important result: the approximation $\mathcal{C}(\delta, \Delta)$ is ε -close to $W^u(x_0)$ and does not ‘jump over’ parts of the manifold, provided the prespecified accuracy parameters were chosen small enough. We stress again that a priori error bounds are not available. Therefore, one always needs to maintain a certain amount of scepticism about a given manifold computation and check the accuracy recomputing the manifold with increased accuracy and comparing the results.

6. Conclusions. We described a general algorithm for computing the global k -dimensional stable or unstable manifold of an equilibrium in an n -dimensional vector field, for any $1 < k < n$. This method can also be used to compute the stable or unstable manifold of an invariant m -torus, provided this m -torus and an approximation of its linear stable and unstable directions can be found; for a method to compute the latter we refer to [3] and also [25, 26].

Our algorithm approximates a first piece of the unstable manifold up to a given geodesic distance from the equilibrium (or invariant torus) with a given error that goes to zero as the tolerance parameters of the algorithm go to zero. Hence, the approximants are upper semicontinuous to $W^u(x_0)$. Furthermore, as a result of the smoothness of the manifold, the algorithm does not jump over parts of $W^u(x_0)$.

The algorithm has been implemented for the case $k = 2$ and arbitrary n and it was illustrated with a number of examples, including some of vector fields arising in applications. We note that the case $n \geq 4$ requires visualizing a two-dimensional object in a high-dimensional space; see Section 4.2.

An implementation of the case $k = 3$ is of much greater complexity and remains work for the future. The unstable manifold is then approximated by a set of 2-spheres, which are represented by simplicial complexes C_i consisting of triangles. The shells between neighboring 2-spheres consist of tetrahedra, which results in more complex overall data management. Moreover, one encounters the general problem of visualizing a three-dimensional object. This makes it difficult to communicate the results, even though some packages, for example Geomview [28], allow for inspection of three-dimensional manifold data.

It seems a fair assessment that in the foreseeable future growth algorithms for global manifolds will be restricted to the case $k = 2$ of two-dimensional manifolds.

In fact, even algorithms for $k = 2$ are largely still in development and have not been used widely. As a consequence, we believe that computing two-dimensional manifolds in problems arising in applications is still in its infancy, with a lot of interesting work to be done.

Acknowledgments. We thank John Guckenheimer, Andrew Stuart and Alexander Vladimirovsky for helpful discussions. B.K. is supported by an EPSRC Advanced Research Fellowship. Both authors gratefully acknowledge hospitality and support of the Control and Dynamical Systems (CDS) Department at the California Institute of Technology, Pasadena where part of this work was done.

REFERENCES

- [1] R.H. ABRAHAM AND C.D. SHAW, *Dynamics — the geometry of behavior, Part three: global behavior*, Aerial Press, Santa Cruz, 1982-1985.
- [2] A. ARNEODO, P. COULLET, E. SPIEGEL, AND C. TRESSER, *Asymptotic chaos*, Physica D **14**(3) (1985), pp. 327–347.
- [3] H.W. BROER, H.M. OSINGA, AND G. VEGTER, *Algorithms for computing normally hyperbolic invariant manifolds*, Z. angew. Math. Phys. **48** (1997), pp. 480–524.
- [4] M. DELLNITZ AND A. HOHMANN, *The computation of unstable manifolds using subdivision and continuation*, In H.W. Broer, S.A. Van Gils, I. Hoveijn, and F. Takens (eds.), *Nonlinear Dynamical Systems and Chaos PNLDE 19* Birkhäuser, 1996, pp. 449–459.
- [5] M. DELLNITZ AND A. HOHMANN, *A subdivision algorithm for the computation of unstable manifolds and global attractors*, Num. Math. **75** (1997), pp. 293–317.
- [6] E. DOEDEL, T. FAIRGRIEVE, B. SANDSTEDE, A. CHAMPNEYS, Y. KUZNETSOV, AND X. WANG, *AUTO 97: Continuation and bifurcation software for ordinary differential equations*, 1997. <http://indy.cs.concordia.ca/auto/main.html>.
- [7] E.J. DOEDEL, Private communications at the IMA, Minneapolis, October 1997.
- [8] J. GUCKENHEIMER AND P. HOLMES, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*, Springer-Verlag, Second edition, 1986.
- [9] J. GUCKENHEIMER AND A. VLADIMIRSKY, *A fast method for approximating invariant manifolds*, preprint 2003 (<http://www.math.cornell.edu/~vlad/papers/InvMfolds/InvMfolds.pdf>).
- [10] J. GUCKENHEIMER AND P. WOLFORD, *Dynamical systems: Some computational problems*, In D. Schlomiuk (ed.), *Bifurcations and Periodic Orbits of Vector Fields* Kluwer Academic Publishers, 1993, pp. 241–277.
- [11] J. HAUSER AND H.M. OSINGA, *On the geometry of optimal control: the inverted pendulum example*, Proc. American Conference on Control Proceedings of the American Control Conference, Arlington VA, USA, 2001, Volume 2, pp. 1721-1726.
- [12] M.E. HENDERSON *Multiple Parameter Continuation: Computing Implicitly Defined k-manifolds*, IJBC v12(3), pages 451-76 *Int. J. Bif. Chaos* **12**(3) (2002), pp. 451-476.
- [13] M.E. HENDERSON, *Computing invariant manifolds by integrating fattened trajectories*, preprint 2003
- [14] D. HOBSON, *An efficient method for computing invariant manifolds of planar maps*, J. Comput. Phys. **104**(1) (1993), pp. 14–22.
- [15] J.G. HOCKING AND G.S. YOUNG, *Topology*, Dover Publications, 1988.
- [16] A. JADBABAIE, J. YU, AND J. HAUSER, *Unconstrained receding horizon control: stability and region of attraction results*, IEEE Proc. Conf. Decision & Control 1999.
- [17] M.E. JOHNSON, M.S. JOLLY, AND I.G. KEVREKIDIS, *Two-dimensional invariant manifolds and global bifurcations: some approximation and visualization studies*, Numerical Algorithms **14** 1997, pp. 125–140.
- [18] M.E. JOHNSON, M.S. JOLLY, AND I.G. KEVREKIDIS, *The Oseberg transition: visualization of global bifurcations for the Kuramoto-Sivashinsky equation*, Int. J. Bif. Chaos **11**(1) (2001), pp. 1–18.
- [19] E.N. LORENZ, *Deterministic nonperiodic flows*, J. Atmospheric Sci. **20** (1963), pp. 130–141.
- [20] B. KRAUSKOPF AND H.M. OSINGA, *Globalizing two-dimensional unstable manifolds of maps*, Int. J. Bif. Chaos **8**(3) (1998), pp. 483–503.
- [21] B. KRAUSKOPF AND H.M. OSINGA, *Growing 1D and quasi 2D unstable manifolds of maps*, J. Comp. Phys. **146**(1) (1998), pp. 404–419.
- [22] B. KRAUSKOPF AND H.M. OSINGA, *Two-dimensional global manifolds of vector fields*, CHAOS **9**(3) (1999), pp. 768–774.

- [23] B. KRAUSKOPF AND H.M. OSINGA, *Visualizing the structure of chaos in the Lorenz system*, Computers and Graphics **26**(5) (2002), pp. 815–823.
- [24] YU. KUZNETSOV, *Elements of Applied Bifurcation Theory*, Springer-Verlag, Berlin, 1995.
- [25] H.M. OSINGA, *Non-orientable manifolds of periodic orbits*, in B. Fiedler, K. Gröger, and J. Sprekels (Eds.) Proc. Int. Conf. Diff. Eqs. **2** World Scientific, Singapore, 2000, pp. 922–924.
- [26] H.M. OSINGA, *Non-orientable manifolds in three-dimensional vector fields*, Int. J. Bif. Chaos **13**(3) (2003), pp. 553–570.
- [27] J. PALIS AND W. DE MELO, *Geometric Theory of Dynamical Systems*, Springer-Verlag, Berlin, 1982.
- [28] M. PHILLIPS, S. LEVY, AND T. MUNZNER, *Geomview: An Interactive Geometry Viewer*, Notices of the Amer. Math. Soc. **40** (1993), pp. 985–988; software and the accompanying manual are available at <http://www.geom.umn.edu/>.
- [29] W.C. RHEINBOLDT, *Methods for Solving Systems of Nonlinear Equations* (Second Edition). CBMS-NSF Regional Conference Series in Applied Mathematics **70**, SIAM Publications, Philadelphia, 1998.
- [30] S.H. STROGATZ, *Nonlinear Dynamics and Chaos*, Addison Wesley, 1994.
- [31] A.M. STUART, AND A.R. HUMPHRIES, *Dynamical Systems and Numerical Analysis*, Cambridge University Press, 1996.
- [32] M. SPIVAK, *Differential Geometry*, Second edition, Publish or Perish, 1979.
- [33] A.J. VAN DER SCHAFT, *L_2 -Gain and Passivity Techniques in Nonlinear Control*, Lecture Notes in Control and Information Sciences **218**, Springer-Verlag, Berlin, 1994.
- [34] P. WORFOLK, Private communications at the Geometry Center, Minneapolis, May 1997.