RC23919 (W0603-029) March 3, 2006
Mathematics

IBM Research Report

Covering an Invariant Manifold with Fat Trajectories

Michael E. Henderson
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Covering an Invariant Manifold with Fat Trajectories

Michael E. Henderson

IBM Research Division, T. J. Watson Research Center, Yorktown Heights, NY 10598, USA. mhender@watson.ibm.com,
www.research.ibm.com/people/h/henderson/

Summary. Invariant manifolds are important objects in the study of dynamical systems, as well as several applications. They are challenging to compute because even in simple systems they can be very complicated surfaces, demanding adaptive schemes to deal with large curvatures.

This paper describes a method that represents the invariant manifold as a set of circular disks in the tangent space, projected onto the manifold which overlap and cover the manifold. These disks are found by integrating fat trajectories, which add tangent and curvature information to the usual point in phase space, and integrates these quantities along a trajectory.

Using a covering eliminates the usual problems with advancing front approaches, and the dual of the covering is a triangulation, should one be needed.

1 Introduction

One of the more important concepts in dynamical systems is that of an invariant manifold. By dynamical system we mean a flow in a phase space $\mathbb{R}^n$

$$\frac{\mathrm{d}x}{\mathrm{d}t} = f(x), \qquad x \in \mathbb{R}^n \tag{1}$$

An *invariant set* of points in phase space is such that points on the trajectory $t \in (-\infty, \infty)$ passing through any point in the set are also in the set. So a collection of trajectories through a discrete set of initial points is an invariant set. If the initial points lie on a smooth curve and the flow $f(x)$ is smooth, the invariant set will be a smooth surface. This is a consequence of the smooth dependance of trajectories on initial conditions (Figure 1.)

Any point in the initial set can be moved forward or backward along a trajectory without changing the invariant manifold, so the curve of initial points defining an invariant manifold is not unique, and a smooth invariant manifold need not be defined by a smooth curve. If a smooth curve of initial points M_0 can be found for an invariant manifold M, M_0 is called a *global transversal* of M.

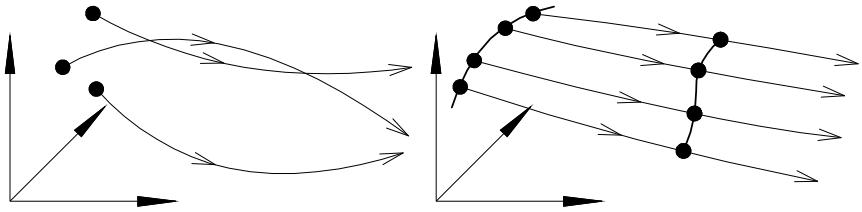

Fig. 1. Invariant sets. (left) defined by a set of points, (right) defined by points on a smooth curve.

For a large class of invariant manifolds M_0 is part of the definition of the manifold. For example, the unstable invariant manifold of a hyperbolic fixed point is the image of a small ball in the unstable eigenspace of the fixed point. However, there are interesting invariant manifolds for which finding a global transversal is part of the problem. A periodic orbit can be found by finding a fixed point of a Poincar'e return map. The fixed point is a global transversal. Not all invariant tori have a closed global transversal, but a torus which contains a quasiperiodic motion does, and a global transversal can be found which is an invariant circle of a return map that is similar to the Poincaré return map [18], [20]. There are better ways to compute periodic orbits [3] and quasiperiodic tori [17], [19], which find M_0 and M together by solving a larger nonlinear system. The literature on all these problems is extensive, and the citations above are not meant to be exhaustive.

Some commonly computed invariant manifolds are summarized in Table 1.

Motion	Geometry	M_0
Fixed Pt.	Point	Point
Periodic Motion	Closed Curve	Point
Heteroclinic Motion	Curve connecting two Fixed Pts.	Point
Quasiperiodic Motion	Torus	Closed curve
Unstable manifold of hyperbolic equilibrium	$\mathbb{R} \times (k-1)$ Sphere	$(k-1)$ Sphere
Inertial manifold	Attracting kd manifold	$(k-1)$d manifold

Table 1. Some common invariant manifolds and the manifolds of starting points which define them.

Certain complex behaviors in dynamical systems are associated with particular configurations of invariant manifolds. However, they can also be useful

by themselves. In orbital mechanics, for instance, it has been proposed [5], [10] to use the unstable manifold of periodic motions (orbits) to design trajectories for spacecraft "missions". These trajectories start in an unstable orbit about a planet or Lagrange point. To stay in the unstable orbit small thrusts are needed. By choosing an appropriate point and direction on the unstable manifold and burning a small amount of fuel, the vehicle can coast along a trajectory on the invariant manifold, and reach certain destinations with no further expense of fuel. The destination might be another unstable periodic orbit, which would allow the spacecraft to return home.

Invariant manifolds are also commonly used in fluid flow visualization, where they are called stream surfaces. In experiments, smoke or dye is introduced into a steady flow along a wire or a tube, and swept downstream. Flow visualization software simulates the experiment by computing the image of the wire under the flow.

In this paper we describe an algorithm for computing a well distributed set of points on a two dimensional invariant manifold when a global transversal is given (a curve in $\mathbb{R}^n$). The algorithm is described in detail for invariant manifolds of dimension two and greater in [9]. The points are spaced along a set of trajectories, and the trajectories are spaced by "fattening" the trajectories. That is, trajectories are not allowed to pass into an interval around the other trajectories.

2 Basic definitions

The "forward" part of an invariant manifold $M^+(M_0)$ consists of all trajectories which start at a point on a smooth curve $M_0 \subset \mathbb{R}^n$. There is also a "backward" part $M^-(M_0)$, found by integrating trajectories backward in time from M_0. This is simply a change of the sign of $f(x)$, so in what follows we drop the superscripts $\pm$, and consider only the forward image of M_0.

The "natural" parameterization of M uses the coordinate σ of a parameterization of M_0, and the time t. Any point on an invariant manifold $M(M_0)$ can be written as $x(\sigma, t) \in \mathbb{R}^n$ where

$$M(M_0) = \left\{ x(\sigma,t) \;\middle|\; x(\sigma,0) = M_0(\sigma), \quad \frac{\mathrm{d}}{\mathrm{d}t} x(\sigma,t) = f(x(\sigma,t)) \right\}.$$

In many interesting cases the natural parameterization is poor (Figure 2). The tangent vectors of the coordinate lines of the natural parameterization are $x^i_{,\sigma}(\sigma,t)$, and $x^i_{,t}(\sigma,t) = f^i(x(\sigma,t))$. A poor parameterization is one where these become nearly linearly dependant, and/or become large or small in norm.

We use the usual tensor notation [16], where the superscript refers to the coordinates of a vector $x \in \mathbb{R}^n$. The subscript with a comma refers to the derivative with respect to the subscript. We will also use the Einstein

summation convention, where the appearance of an index twice in a product indicates a sum over that index. The inner product of two vectors is written $x^p y^p$, to mean $\sum_p x^p y^p$. We will try to use $p, q, r, ...$ for indices which are summed over, and $i, j, ...$ for other indices.

The metric g at a point on $M(M_0)$ is the 2×2 matrix

$$g = \begin{pmatrix} x^p_{,\sigma} x^p_{,\sigma} & x^p_{,\sigma} x^p_{,t} \\ x^p_{,t} x^p_{,\sigma} & x^p_{,t} x^p_{,t} \end{pmatrix}.$$

A "good" parameterization is one for which g is everywhere close to the identity. That is $x_{,\sigma}$ and $x_{,t}$ are unit vectors, and $x_{,t}$ is orthogonal to $x_{,\sigma}$.

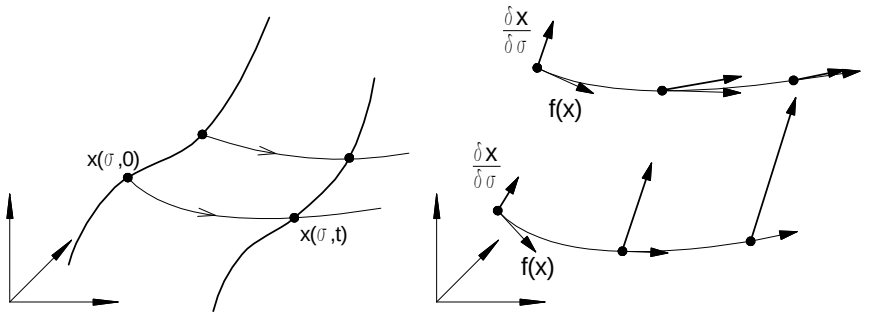

Fig. 2. The natural parameterization. (left) A "good" parameterization, (right) two types of "poor" parameterizations. The lower trajectory has tangent vectors that are roughly orthogonal (that is, g is diagonal), but not unit vectors. The upper trajectory suffers from shear, where the tangent vectors are far from orthogonal (g has large off diagonal elements.

One way to understand the literature on computing invariant manifolds is to consider how the natural parameterization M is improved. (A recent survey [14] describes and constrasts the various approaches.) [11] and [13] use a diagonal scaling, while [6], [7] and [15] use an upper triangular scaling. These scalings are done indirectly, by adapting a mesh, and if the coordinate curves no longer align with the trajectory, some sort of interpolation must be done. The approach described here uses a parameterization that is locally Euclidean near a trajectory (i.e. g in the new parameterization is the identity), and advances points, tangents and curvature along trajectories (a *fat* trajectory).

3 Fat Trajectories

In order to build an interval about a trajectory we need the tangent space and curvature of the invariant manifold M. The interval trajectory will be the set

of points which when projected to the tangent space of the nearest point on the fat trajectory, lie inside a disk about the origin. The radius of the disk will be allowed to vary along the trajectory according to the curvature of the trajectory. A good choice is $R = \sqrt{\epsilon/2 \parallel x'' \parallel}$, where x'' is the curvature and ϵ controls the distance between the tangent space and the manifold (see [8] for details).

In the natrual parameterization the t tangent vector is $x_{,t} = f(x)$. The σ tangent must be integrated along a trajectory starting at $M_0(\sigma)$ –

$$\frac{\mathrm{d}}{\mathrm{d}t} x^i_{,\sigma} = f^i_{,p} x^p_{,\sigma}. \tag{2}$$

With a little differential geometry it can be shown [9] that an orthonormal basis x^i_0, x^i_1 for the tangent space which changes as little as possible along a trajectory (Figure 3) evolves according to

$$\frac{\mathrm{d}}{\mathrm{d}t} x^i_{,j} = f^i_{,p} x^p_{,j} - \left(x^p_{,r} f^p_{,q} x^q_{,j} \right) x^i_{,r} \tag{3}$$

Equation 3 is similar in form to Equation 2, but a linear combination of the tangent vectors has been subtracted, and this maintains the orthonormality of the basis. Evolution equations can also be found for the curvature (or more precisely the derivatives of the tangent vectors). In the natural parameters

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} x^i_{,t,t} &= f^i_{,p} x^p_{,t} \\ \frac{\mathrm{d}}{\mathrm{d}t} x^i_{,\sigma,t} &= f^i_{,p} x^p_{,\sigma} \\ \frac{\mathrm{d}}{\mathrm{d}t} x^i_{,\sigma,\sigma} &= f^i_{,p} x^p_{,\sigma,\sigma} + f^i_{,p,q} x^p_{,\sigma} x^q_{,\sigma} \end{aligned} \tag{4}$$

and in the orthonormal basis

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t} x^i_{,j,k} = {} & f^i_{,p} x^p_{,j,k} + f^i_{,p,q} x^p_{,j} x^q_{,k} \\ & -(x^p_{,r} f^p_{,q} x^q_{,j}) x^i_{,r,k} - (x^p_{,r} f^p_{,q} x^q_{,k}) x^i_{,r,j} \\ & -(x^p_{,w} f^p_{,q} x^q_{,j,k} + x^p_{,w} f^p_{,q,r} x^q_{,j} x^r_{,k} + x^p_{,j,k} f^p_{,q} x^q_{,w}) x^i_{,w} \end{aligned} \tag{5}$$

Though the expressions are of course more complicated, the form of Equation 5 is the same as Equation 4 except that this time linear combinations of both the second derivative vectors and the tangents have been subtracted.

Initial conditions for the basis at points on M_0 can be found using Gram–Schmidt orthogonalization starting with the natural parameterization. Since the second derivatives in the natural parameterization are easily found they can be transformed into second derivatives in the new basis [9].

This coordinate system is analogous to the Riemannian Normal Coordinates (RNC) used in general relativity to find an inertial frame along geodesics. Here the trajectory plays the role of the geodesic. The coordinate system is also a parallel transport.

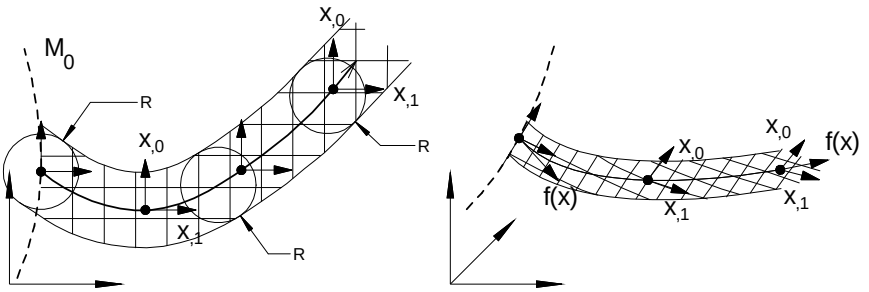

Fig. 3. A sketch of the new coordinate system near a trajectory. (left) Looking "down" on M. (right) the same in space. Note that the flow direction $f(x)$ is not one of the two basis vectors. The lines paralleling the trajectory are a neighborhood on M of the trajectory with width R.

Fat trajectories are neighborhoods of width $R(x)$ about a trajectory, with $R(x)$ varying along the trajectory (Figure 3). To cover M, a set of points is distributed on M_0 using $R(x)$, and fat trajectories are integrated forward from these points. The integration is stopped if the trajectory enters a previously integrated fat trajectory. This may leave uncovered parts of M if the flow expands (which is common). To cover the rest of M we must locate points and construct initial conditions for starting more fat trajectories. To do this we use circular disks in the tangent space of the fat trajectory at points spaced on the trajectory according to $R(x)$. This allows us to use a representation of the boundary of the covered part of M to locate an interpolation point.

4 Flying disks

In [8] the author developed a method of representing manifolds as the union of overlapping spherical balls of different radii. The representation was used to compute implicitly defined manifolds (i.e. solutions of $F(x) = 0$ with $F : \mathbb{R}^n \to \mathbb{R}^{n-k}$), and has been used for computing other types of manifolds as well. The approach computes an approximate "restricted Laguerre–Voronoi" tessellation of M based on the spherical balls. This is instead of the more usual grid, or triangulation, though the dual Delaunay triangulation can be used if a triangulation is required (Figure 4.) Voronoi and Delaunay diagrams are described in [2]. The restricted Laguerre–Voronoi diagram, or restricted

power diagram, is decribed in [1], [4] and [12]. Using the Voronoi tessellation avoids the well known problems with advancing triangulations, and the radius of the spherical ball provides a way of equi-distributing points on M. Roughly, points are no closer than R, or further apart than $2R$.

Below we describe the two dimensional case, but the same approach works in higher dimensions, with a polyhedral Voronoi tessellation instead of polygonal tesselation.

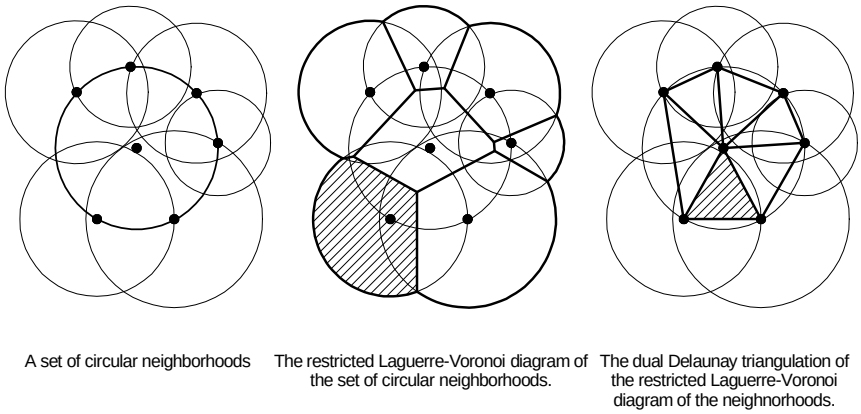

Fig. 4. A set of circular neighborhoods and the corresponding restricted Laguerre-Voronoi and dual Delaunay diagrams.

A triangulation would probably be the first choice to represent a manifold. To iteratively find a set of points on a manifold a point on the boundary would be identified (easily done for a triangulation) and advanced some distance normal to the boundary. The new point would then be used to define a triangle (or simplex in higher dimensions) which is added to the mesh. This keeps the triangulation moving "outward" from the initial point, but there are many cases in which the new triangle is incompatible with the existing triangulation. That is, the new triangle overlaps the existing triangulation.

A covering is a set of neighborhoods centered at points on the manifold. The neighboroods are allowed to overlap, as long as every point on M lies in some neigborhood. A covering does not have the difficulty with compatibility of new neighborhoods as triangulations (they are meant to overlap). However, it is not obvious how to find a point near the boundary of a union of neighborhoods. The polygonal Voronoi tiles provide a way to find a point on the boundary.

Finding the Voronoi tiles is simple. The points are found at the same time, so this is not the usual incremental computation of a Voronoi diagram, where the points are givem. If we have one circular disk, and a square which contains

it, then the boundary of the disk is the part of the circle inside the square. When a second disk is added, the intersection of the circles bounding the two disks lies on a line (in higher dimensions a plane) orthogonal to the line between the centers of the two disks, and the part of the circle on the boundary of the union is the part on the appropriate side of this line. If complementary halfplanes are removed from the squares surrounding the two disks, the part of the boundary of each disk is the part inside the resulting polygon (figure 5). By identifying neighboring disks when a new disk is added, the Voronoi tiles (the part of the disk inside the polygon) are updated by removing complementary halfplanes from the new disk and each disk which it overlaps.

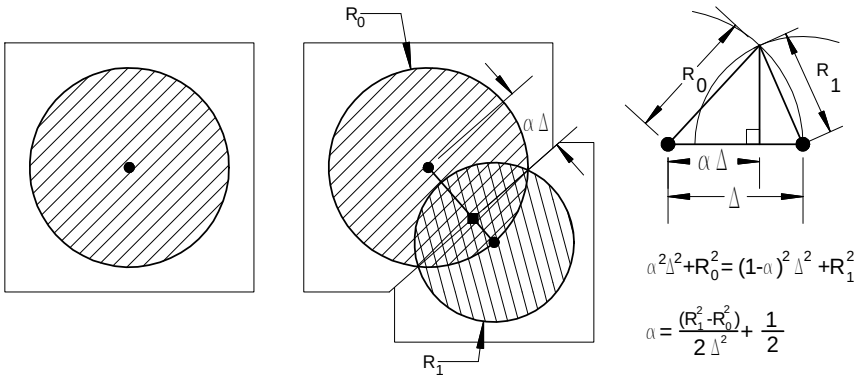

Fig. 5. Updating Voronoi tiles. Each tiles starts as a square, then for each disk which overlaps the disk the polygon is clipped against the line by the intersection of the circles of the two disks.

When disks are in different tangent spaces, they must be projected to a common tangent space before updating the polygons. If the radius of the disk is small relative to the curvature of M, the projection of one disk to the tangent space of an overlapping disk will almost be a circular disk, and the previous procedure can be used to update the polygons. There is an error committed, but the effect is that points that are identified as boundary points may be slightly inside the boundary (Figure 6).

The invariant manifold M is represented as the union of the projections of a set of circular disks onto M. This is a list, or "atlas" of "charts", which consist of a point on M, tangent vectors of M at that point, and a radius (these represent the circular disk), together with a polygon (the Voronoi tile). As points are added to the list, the polygons are updated by clipping the polygon against a line.

To approximate a fat trajectory we start with a point $x_0 \in M_0$, or an interpolated point, and compute the initial orthonormal basis and second derivatives. This forms the first chart on the fat trajectory. The trajectory,

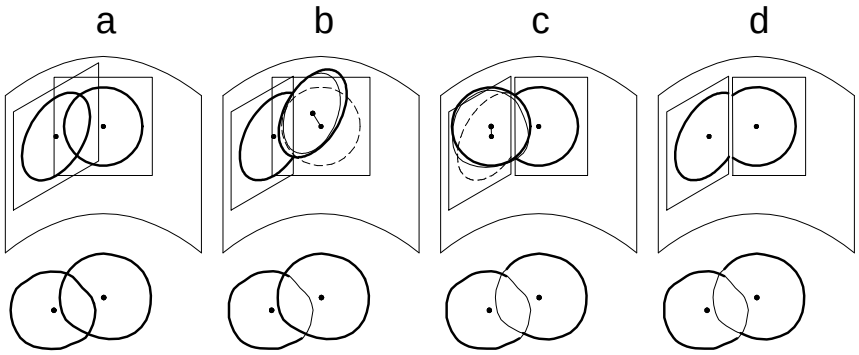

Fig. 6. On a curved surface (a) the update of a disk's polygon is done in the disk's tangent space. The center of each overlapping disk is projected into the tangent space (b), and a circle with the radius of the overlapping disk is used to update the polygon. There is an error involved, since the neighborhood is actually on the manifold (sketched below the surface), and the circle is distorted by the projection onto M and then the projection into the tangent space. This process is then repeated (c) to update the polygon of the neighboring disk. The result (d) is still part of the boundary of the projection onto M, but we may think that a point is on the boundary when it is actually a little inside. The size of this error is of the order of the distance between the tangent space and M on the circle.

tangent space and curvature are integrated a distance R, and another chart is added. This process is repeated until a maximum time is reached, or the trajectory enters an existing chart. A sketch of a fat trajectory that has been covered this way is shown in figure 7.

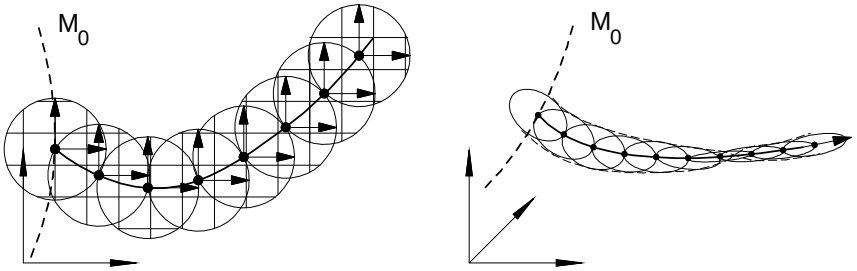

Fig. 7. Circular neighborhoods (charts) along a trajectory. (left) Looking "down" on M. (right) the same in space, showing how the disk "rolls" and "pitches", but does not "yaw".

5 Interpolation

When there are no more points on M_0 which are outside the union of charts (Figure 8), a starting point must be found that is not on M_0. A point in the interior of the union will be inside a fat trajectory, so we use the polygons to find a point near the boundary, where the new trajectory will leave the interior of the union.

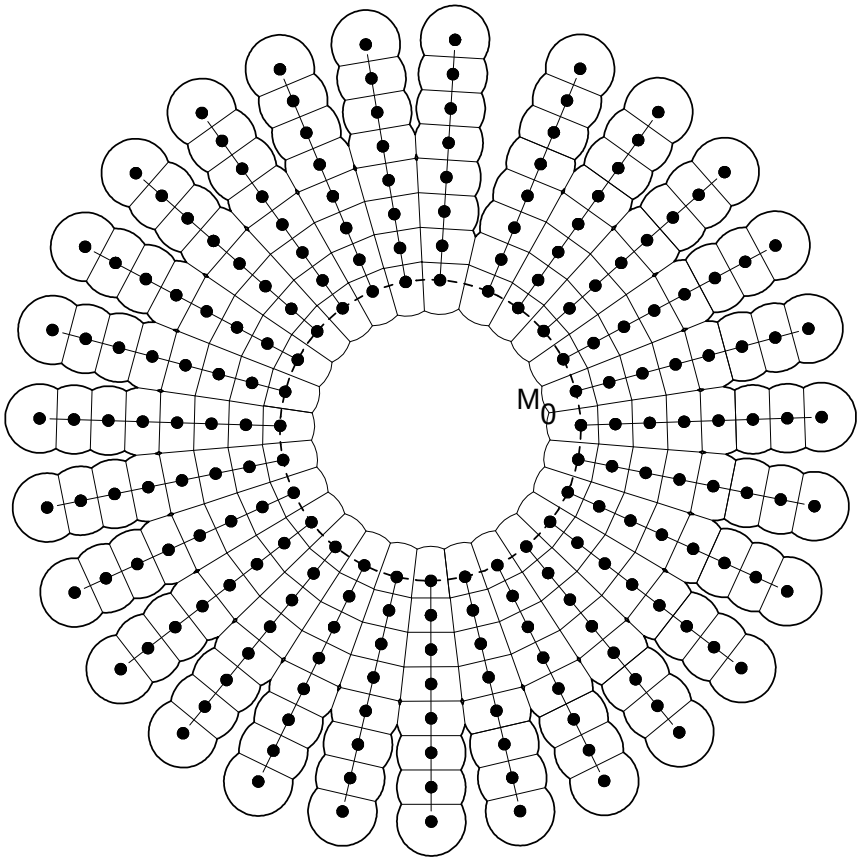

Fig. 8. When the manifold of starting points M_0 is covered, some other point must be found that can be used to start a new trajectory.

In [9] the author used an argument based on a modified nonlinear optimization problem to show that such a point exists. There is a technical requirement that is satisfied for 2d surfaces once M_0 is covered, and the variation in $f(x)$ over the disk must be small relative to the radius of the disk. The optimization problem looks for a point on the boundary of the union which is furthest back

(locally) toward M_0 along trajectories on M. While there is no objective in this optimization problem, the usual optimality conditions must hold. These conditions are that an optimum be a stable fixed point of a modified flow (on a boundary the component of $f(x)$ normal to the boundary is projected out). The problem is posed on the part of M which is *outside* the disks, and not further than a maximum time T_{max} from M_0 (measured along trajectories).

There can be no fixed points on the extrior of the union, since it takes an infinite amount of time to reach the fixed point, and a maximum time has been imposed. The point furthest back toward M_0 must therefore lie on the boundary of the disks. For a point on the boundary to be a fixed point, the flow vector must lie in the positive cone of normal vectors. This is just another way of introducing Lagrange multipliers. There are only two types of point on the boundary, those which lie on a single circle, and those at the intersection of two circles. For a fixed point on a single circle $f(x)$ must be parallel to the normal of the circle, and point away from the center of the circle. That is, *the extension of the flow vector on the boundary backward in time passes through the center of the circle* (Figure 9). However, such a point cannot be a minimum, since it lies on a circle which curves backward in time.

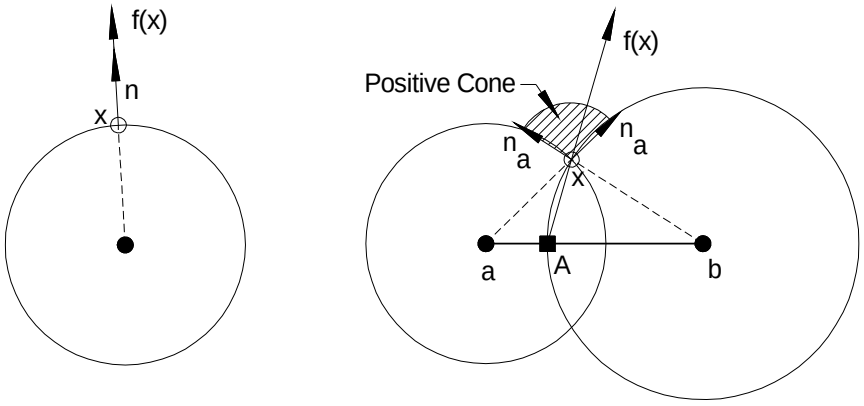

Fig. 9. (left) A point x on the boundary of a single disk is a fixed pointof the modified flow if $f(x)$ is parallel to the normal and points in the "outward" direction. That is, the extension of $f(x)$ backward in time passes through the center of the disk. Moving x a little in either direction on the circle moves x further back towards M_0, so it is not the "minimum". (right) A point x at the intersection of two circles is a fixed point of the modifed flow if $f(x)$ lies in the positive cone formed from the two normals. That is, the extension of $f(x)$ backward in time crosses the line between centers a and b (point A). This is a local "minimum", and a trajectory started at A – if f does not change too rapidly over the radius of a disk – will pass out of the interior of the disks near x. Initial values for the tangents and curvature can be interpolated from centers a and b.

At fixed points of the modified flow lying at the intersection of two circles must have $f(x)$ in the positive cone formed by the normals of the two circles. The normals are parallel to lines starting at the center and passing through the intersection point (Figure 10). That is, *the extension the flow vector at the boundary point backward in time crosses the interior of the edge between the two centers.* This point is the minimum, and if we use the point on the edge between the centers to start a new trajectory, the initial values can be interpolated from the values at the two centers, and if $f(x)$ does not vary much over a disk the new trajectory will leave the union near the intersection point.

Intersection points can be easily found from the polygons associated with the disks (Figure 10). They are points where an edge of the polygon crosses the circular boundary of the disk. A list of the disks on the boundary can be maintained (boundary disks have polygons with at least one exterior vertex). To find an interpolation point this list is transversed, and the edges of the polygon are tested for crossing. If one endpoint is inside and the other outside this is trivial. If both endpoints are outside, the distance between the edge and the center being less than the radius indicates a crossing.

6 Example

As an illustration, we consider a periodically forced pendulum with damping (this example is from [21]). When the forcing is zero, the phase space consists of a set of hyperbolic fixed points at $x_1 = (2n+1)\pi$, $(x_1)_{,t} = 0$, where the pendulum points straight up, and centers at $x_1 = 2n\pi$, $(x_1)_{,t} = 0$ with the pendulum pointing down. With periodicity the phase space can be reduced to $x_1 \in [-\pi, \pi]$. Figure 11 shows the unperturbed nonlinear single pendulum. Gravity acts on the pendulum bob and the equations are

$$\frac{\mathrm{d}}{\mathrm{d}t} x_1 = x_2$$

$$\frac{\mathrm{d}}{\mathrm{d}t} x_2 = -\sin x_1$$

Without the forcing and dmping there is an energy $E = x_2^2/2 - \cos x_1$ that is conserved on trajectories. For initial energies $E > 1$ the pendulum "runs", that is x_1 continually increasing or decreasing depending on the initial velocity. For $E < 1$ the pendulum oscillates about the downward pointing fixed point. If the pendulum is started with $E = 1$ it will swing to the top and stop. This last is a heteroclinic orbits shown as dark curve in Figure 11 center.

The perturbated equations, analyzed in [21] are

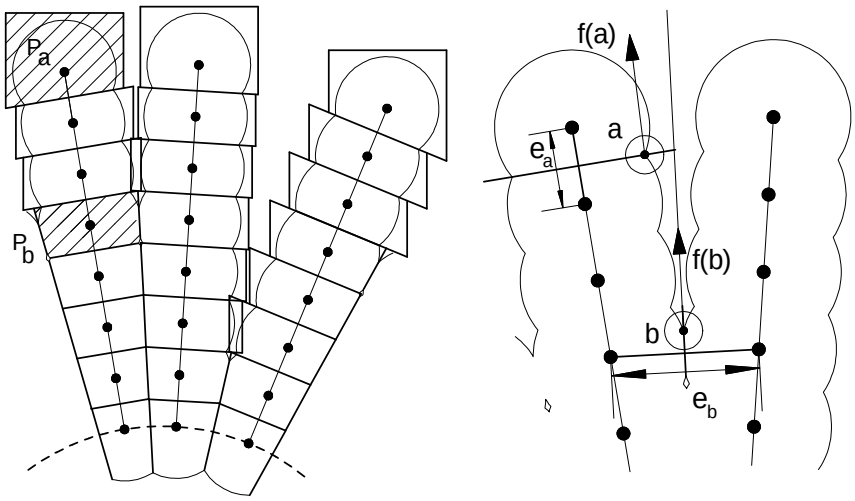

Fig. 10. Interpolating: (left) Some of the disks and polygons from Fig. 8. Polygons with a vertex outside the disk indicate that the disk is on the boundary. We highlight two polygons, P_a and P_b. (right) The edges of P_a and P_b which cross the boundary of the disks, and two of the crossing points, a and b. All of the other edges of the polygons have been removed. The centers that these two edges separate form an edge. At point a the flow vector $f(a)$ extended backward does not intersect e_a, while at the point b $f(b)$ extends backward to cross e_b. A trajectory started at the point where the two cross will leave the union of the disks, and initial values can be interpolated between the centers at the ends of e_b.

$$\frac{\mathrm{d}}{\mathrm{d}t}x_1 = x_2$$

$$\frac{\mathrm{d}}{\mathrm{d}t}x_2 = -\sin x_1 + \epsilon\left(\gamma \sin(\Omega x_3)\sin x_1 - \delta x_2\right))$$

$$\frac{\mathrm{d}}{\mathrm{d}t}x_3 = 1$$

The perturbation is time dependant, so time is introduced as a phase space coordinate to make the flow autonomous (a standard trick called *suspending* the flow).

For small perturbations ($\epsilon << 1$) there is a critical forcing amplitude

$$\gamma^* = \frac{4\delta}{\pi\Omega}\sinh\frac{\pi\Omega}{2}$$

For $\gamma < \gamma^*$ the damping removes more "energy" than the periodic forcing puts into the system, and the pendulum eventually will spiral into the fixed

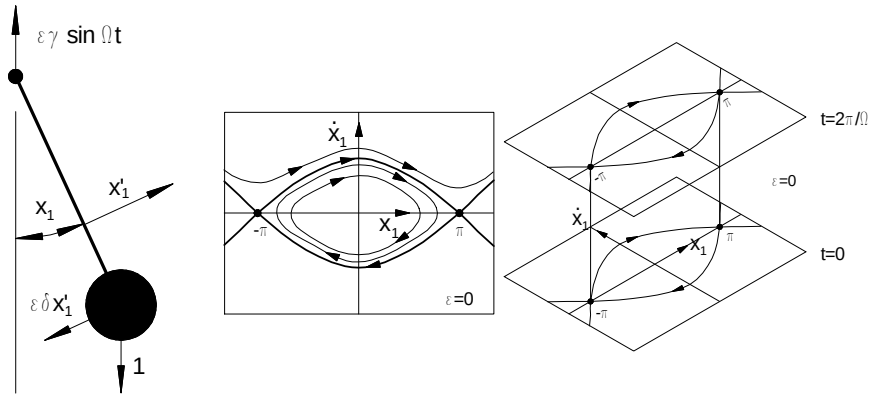

Fig. 11. The periodically forced pendulum from [21]. (left) A periodic vertical force is applied to the pivot, as well as a damping. (center) the behavior of the pendulum when $\epsilon = 0$. If the energy $x_1^2/2 - \cos x_1$ is greater than one the pendulum swings around and around. If the energy is less than one the pendulum oscillates (there is no damping at $\epsilon = 0$. When the energy is exactly one, the pendulum comes to rest with the bob above the pivot. (right) with $\epsilon > 0$ time becomes a variable, periodic over $2\pi/\Omega$. Following [21] we will use this box to display the image of a line of initial points.

point at zero (which is now the straight line $(0, 0, t)$). For $\gamma \geq \gamma^*$ heteroclinic tangles appear (a type of chaotic motion).

For illustration we chose M_0 to be the line $(x_1, x_2, x_3) = (3.0, -0.1, x_3)$, which is near one of the unstable fixed points for $\epsilon = 0$. Figure 12 shows the surface that was computed. The time coordinate x_3 is periodic with period $2\pi/\Omega$, and Figure 12 shows sixteen periods of x_3. If we use the same computational results and collapse it to two periods of x_3 we can see some of the structure that leads to a heteroclinic tangle and chaotic motion.

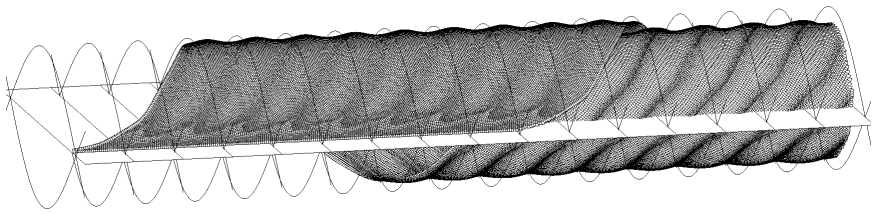

Fig. 12. The periodically forced pendulum. $\epsilon = .2$, $\gamma = 1.5$, $\delta = .2$, and $\Omega = 5$. The calculation used 176 points on seven replicas of the fundemental region $[0, \pi/\Omega)$, and those 176 start points created 84,943 disks. In addition, 48 interpolations were needed, for a total of 91,240 disks.

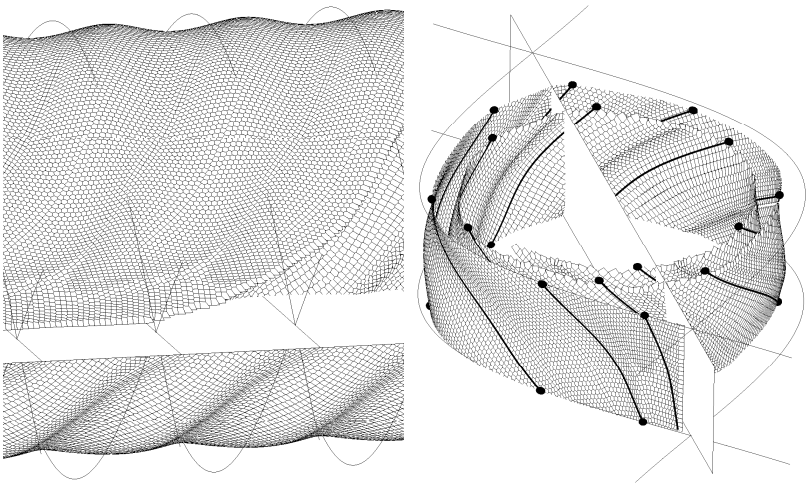

Fig. 13. (left) A closer view of the invariant manifold in Figure 12. (right) The same invariant manifold but brought back to a single period in the forcing. The dark black line is a trajectory starting near the fixed point at $x_1 = \pi$. The mapping from the plane $t = 0$ to $t = 2\pi/\Omega$ is used in the analysis, and the black dots are the orbit of one point under that map. For these parameters the motion decays to the fixed point of the map at $(0, 0)$

References

1. F. Aurenhammer, *Power diagrams: Properties, algorithms and applications*, SIAM Journal of Computing, 16 (1987), pp. 78–96.
2. ———, *Voronoi diagrams – a survey of a fundamental geometric data structure*, ACM Computing Surveys, 23 (1991), pp. 345–405.
3. E. J. Doedel, *Nonlinear numerics*, International Journal of Bifurcation and Chaos, 7 (1997), pp. 2127–2143.
4. H. Edelsbrunner, *The union of balls and its dual shape*, Discrete and Computational Geometry, 13 (1995), pp. 415–440.
5. G. Gomez, W. S. Koon, M. W. Lo, J. E. Marsden, J. Masdemont, and S. D. Ross, *Invariant manifolds, the spatial three-body problem and space mission design*, in International Conference on Differential Equations, Berlin, 1999, pp. 1167–1181.
6. J. Guckenheimer and A. Vladimirsky, *A fast method for approximating invariant manifolds*, SIAM J. Appl. Dyn. Systems, 3 (2004), pp. 232–260.
7. J. Guckenheimer and P. Worfolk, *Dynamical systems: Some computational problems*, in Bifurcations and Periodic Orbits of Vector Fields, D. Schlomiuk, ed., Kluwer Academic Publishers, 1993, pp. 241–277.
8. M. E. Henderson, *Multiple parameter continuation: Computing implicitly defined k–manifolds*, International Journal of Bifurcation and Chaos, 12 (2002), pp. 451–476.

9. ———, *Computing invariant manifolds by integrating fat trajectories*, SIAM Journal on Applied Dynamical Systems, 4 (2005), pp. 832–882.
10. K. C. Howell, B. T. Barden, and M. W. Lo, *Application of dynamical systems theory to trajectory design for a libration point mission*, The Journal of the Astronautical Sciences, 45 (1997), pp. 161–178.
11. J. P. M. Hultquist, *Constructing stream surfaces in steady 3D vector fields*, in IEEE Proceedings Visualization '92, Boston, October 1992, pp. 171–178.
12. H. Imai, M. Iri, and K. Murota, *Voronoi diagram in the laguerre geometry and its applications*, SIAM Journal on Computing, 14 (1985), pp. 93–105.
13. 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.
14. B. Krauskopf, H. Osinga, E. Doedel, M. Henderson, J. Guckenheimer, A. Vladimirsky, M. Dellnitz, and O. Junge, *A survey of methods for computing (un)stable manifolds of vector fields*, International Journal of Bifurcation and Chaos, 15 (2005), pp. 2127–2143.
15. B. Krauskopf and H. M. Osinga, *Computing geodesic level sets on global (un)stable manifolds of vector fields*, SIAM Journal on Applied Dynamical Systems, 2 (2003), pp. 546–569.
16. D. Lovelock and H. Rund, *Tensors, Differential Forms, and Variational Principles*, John Wiley & Sons, New York, 1975.
17. B. Rassmussen, *Numerical Methods for the Continuation of Invariant Tori*, PhD thesis, Georgia Institute of Technology, School of Mathematics, December 2003.
18. V. Reichelt, *Computing invariant tori and circles in dynamical systems of fixed points*.
19. F. Schilder, H. M. Osinga, and W. Vogt, *Continuation of quasi-periodic invariant tori*, SIAM J. Applied Dynamical Systems, 4 (2005), pp. 459–488.
20. M. van Veldhuizen, *A new algorithm for the numerical approximation of an invariant curve*, SIAM J. Sci. Stat. Comput., 8 (1987), pp. 951–962.
21. S. Wiggins, *Global Bifurcations and Chaos*, no. 73 in Applied Mathematical Sciences, Springer–Verlag, New York, 1988.