Polygonal Approximation of Flows

Erik Boczko
Department of Biomedical Informatics
Vanderbilt University
Nashville, TN 37235
E-mail: erik.m.boczko@vanderbilt.edu

William D. Kalies*
Department of Mathematical Sciences
Florida Atlantic University
Boca Raton, FL 33431
E-mail: wkalies@fau.edu

Konstantin Mischaikow †
Center for Dynamical Systems and Nonlinear Studies
Georgia Institute of Technology
Atlanta, GA 30332 USA
E-mail: mischaik@math.gatech.edu

May 12, 2006

Abstract

The analysis of the qualitative behavior of flows generated by ordinary differential equations often requires quantitative information beyond numerical simulation which can be difficult to obtain analytically. In this paper we present a computational scheme designed to capture qualitative information using ideas from the Conley index theory. Specifically we design an combinatorial multivalued approximation from a simplicial decomposition of the phase space, which can be used to extract isolating blocks for isolated invariant sets. These isolating blocks can be computed rigorously to provide computer-assisted proofs. We also obtain local conditions on the underlying simplicial approximation that guarantees that the chain recurrent set can be well-approximated.

KEYWORDS: chain recurrence, combinatorial dynamics, Conley index, flows from ODE's

AMS SUBJECT CLASSIFICATION: 37B20, 37B25, 37B30, 37B35, 37C10, 37C70, 37M99

*Research supported in part by NSF grant DMS-9973331.
†Research supported in part by NSF grant DMS-9505116 and DMS-9805584.

1 Introduction

A significant portion of both dynamical systems theory and numerical analysis has its roots in the study of ordinary differential equations. However, due to the historical development of these two disciplines a dichotomy has developed. On the numerical side the primary focus has been on the approximation of individual orbits. In dynamics the central focus is on the structure of invariant sets. In fact, one of the major accomplishments of dynamical systems has been to demonstrate that the global dynamics of ordinary differential equations can only be completely understood through an understanding of the behavior of sets of trajectories rather than individual orbits. On the other hand, the high level of generality at which dynamical systems is conducted makes it difficult to rigorously apply these ideas to particular equations. For this one needs quantitative estimates that are often most easily obtained by numerical approximation.

That the two approaches have different strengths is of course well known. Dynamicists have for a long time used numerical simulations to investigate and demonstrate dynamical properties of differential equations, and likewise, numerical analysts have begun to use the ideas from dynamical systems to obtain a systematic understanding of the implications of applying numerical schemes and algorithms (see [1, 5, 9] and references therein).

However, to the best of our knowledge there are no general computational schemes which are designed explicitly to capture the qualitative dynamics of ordinary differential equations. The purpose of this paper is to introduce such a scheme. Since the aim is very ambitious and extends beyond our current abilities of implementation, we will begin with an introduction describing how Conley's approach to dynamical systems can provide a theoretical basis for developing a unified computational approach to dynamical systems. Indeed, in [6] we showed that Conley's approach to dynamics can be made completely algorithmic for discrete dynamical systems; certain issues remain unresolved for flows, but here we describe a framework in which to address some of these issues and in which to develop algorithms for computations.

To be concrete, throughout this paper we will study a differential equation of the form

$$\dot{x} = f(x), \quad x \in \mathbb{R}^n \tag{1}$$

which generates a flow

$$\varphi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n \tag{2}$$

We are interested in approximating the dynamics on a polyhedron $X \subset \mathbb{R}^n$.

From the dynamical systems point of view the object of interest is the structure of the maximal invariant set in X,

$$\operatorname{Inv}(X, \varphi) := \{x \in X \mid \varphi(\mathbb{R}, x) \subset X\},$$

and our goal is to develop computational methods to capture this information.

The first observation that needs to be made is that, because of local and global bifurcations, invariant sets are not stable with respect to perturbations. Any numerical method introduces errors and hence should be thought of as a perturbation to the system of interest. Therefore, while it is possible to develop numerical methods to find particular orbits, such as fixed points or periodic orbits, or particular structures, such as unstable manifolds or invariant tori, to study general invariant sets we must proceed indirectly.

We begin with the following concept.

Definition 1.1 A compact set $N \subset \mathbb{R}^n$ is an *isolating neighborhood* if $\text{Inv}(N, \varphi) \subset \text{int}\, N$. An invariant set is an *isolated invariant set* if it is the maximal invariant set of an isolating neighborhood.

The method we are developing here is intended to elucidate the structure of isolated invariant sets. To do this we need to be able to decompose invariant sets in a robust manner. This leads to the following idea.

Definition 1.2 Let $\epsilon, \tau > 0$. An (ϵ, τ)-*chain* from x to y in X is a finite sequence $j = 1, \ldots, J$ of pairs

$$\{(z_j, t_j) \subset X \times [\tau, \infty) \mid x = z_1, y = z_J, \|\varphi(t_j, z_j) - z_{j+1}\| < \epsilon, \varphi([0, t_j], z_j) \subset X\}.$$

The (ϵ, τ)-*chain recurrent set* of X is

$$\mathcal{R}_{\epsilon,\tau}(X) := \text{cl}\left(\{x \in X \mid \exists \text{ an } (\epsilon, \tau) - \text{chain from } x \text{ to } x\}\right)$$

The *chain recurrent set* is

$$\mathcal{R}(X) := \bigcap_{\epsilon > 0} \mathcal{R}_{\epsilon,\tau}(X)$$

and as the notation suggests is independent of τ.

In this context Conley's decomposition theorem [2] is as follows.

Theorem 1.3 *Let $\mathcal{R}^j(X)$, $j = 1, 2, \ldots$, denote the connected components of $\mathcal{R}(X)$. Then there exists a continuous function $V : \text{Inv}(X, \varphi) \to [0, 1]$ such that*

1. *if $x \notin \mathcal{R}(X)$ and $t > 0$, then $V(x) > V(\varphi(t, x))$,*
2. *for each $j = 1, 2, \ldots$ there exists $\sigma_j \in [0, 1]$ such that $\mathcal{R}^j \subset V^{-1}(\sigma_j)$.*

With numerical approximations in mind, we are more interested in $\mathcal{R}_{\epsilon,\tau}(X)$ for a fixed $\epsilon, \tau > 0$, than $\mathcal{R}(X)$. As implied in the decomposition theorem, the number of components of the chain recurrent set is at most countable. However, it can be shown that only finitely many components of $\mathcal{R}_{\epsilon,\tau}(X)$ have a nonempty intersection with $\mathcal{R}(X)$ and hence a nonempty invariant set. This leads to the following result.

Theorem 1.4 *Let $\epsilon, \tau > 0$, and let $\mathcal{R}^j_{\epsilon,\tau}(X)$, $j = 1, 2, \ldots, J$, denote the (finitely many) connected components of $\mathcal{R}_{\epsilon,\tau}(X)$ for which $\text{Inv}(\mathcal{R}^j_{\epsilon,\tau}(X), \varphi) \neq \emptyset$. Then,*

$$\text{cl}\left(\mathcal{R}^j_{\epsilon,\tau}(X)\right)$$

is an isolating neighborhood. Let $M(j) := \text{Inv}(\mathcal{R}^j_{\epsilon,\tau}(X), \varphi)$. Then there exists a continuous function $V : \text{Inv}(X, \varphi) \to [0, 1]$ such that

1. *if $x \notin \cup_{j=1}^J M(j)$ and $t > 0$, then $V(x) > V(\varphi(t, x))$,*
2. *for each $j = 1, 2, \ldots, J$ there exists $\sigma_j \in [0, 1]$ such that $M(j) \subset V^{-1}(\sigma_j)$.*

The sets $M(j)$ are called Morse sets and together form a *Morse decomposition* of $\mathrm{Inv}(X, \varphi)$. The function V is called a *Lyapunov function*. Observe that if $x \in \mathrm{Inv}(X, \varphi) \setminus \cup_{j=1}^{J} M(j)$ then there exist two distinct Morse sets M and M' such that

$$\omega(x, \varphi) := \bigcap_{t \geq 0} \mathrm{cl}(\varphi([t, \infty), x) \subset M$$

and

$$\alpha(x, \varphi) := \bigcap_{t \geq 0} \mathrm{cl}(\varphi((-\infty, -t], x) \subset M'$$

With this in mind, it is clear that we can also think of a Morse decomposition of $\mathrm{Inv}(X, \varphi)$ as a finite collection of compact invariant sets indexed by a finite set $\mathcal{J}$, i.e.

$$\mathcal{M}(\mathrm{Inv}(X, \varphi)) = \{M(j) \mid j \in \mathcal{J}\}$$

with the added condition that there one can impose a strict partial order $>$ on the elements of $\mathcal{J}$ with the property that $j > k$ implies that there is no element $x \in X$ such that

$$\omega(x, \varphi) \subset M(j) \quad \text{and} \quad \alpha(x, \varphi) \subset M(k)$$

Such an ordering is called an *admissible order*.

Given a Morse decomposition, it is possible to produce a variety of additional isolated invariant sets, as the following procedure indicates. Given an indexing set $\mathcal{J}$ and a partial order $>$, $I \subset \mathcal{J}$ is an *interval* if $p, q \in I$ and $p > r > q$ implies that $r \in I$. The set of intervals is denoted by $\mathcal{I}(\mathcal{J}, >)$. An interval I is *attracting* if $p \in I$ and $p > q$ implies that $q \in I$. The set of attracting intervals is denoted by $\mathcal{A}(\mathcal{J}, >)$. For $I \in \mathcal{I}(\mathcal{J}, >)$ define

$$M(I) := \left(\bigcup_{p \in I} M(p) \right) \cup \left(\bigcup_{p, q \in I} C(M(p), M(q)) \right)$$

where

$$C(M(p), M(q)) := \{x \mid \omega(x) \subset M(q), \alpha(x) \subset M(p)\}.$$

It is left to the reader to check that $M(I)$ is an isolated invariant set.

As will become clear in the next section, our computational scheme is designed to find a Morse decomposition and a discrete approximation to the associated Lyapunov function. What remains to be discussed is how one uses the approximation to understand the structure of the invariant set. This is done via the Conley index which can be computed in terms of special isolating neighborhoods.

Definition 1.5 An isolating neighborhood N is an *isolating block* if for every $x \in \partial N$, and every $\epsilon > 0$,

$$\varphi((-\epsilon, \epsilon), x) \not\subset N. \tag{3}$$

The *immediate exit set* for N is given by

$$N^- := \{x \in N \mid \forall \epsilon > 0,\ \varphi((0, \epsilon), x) \not\subset N\}.$$

If N is an isolating block, then the *Conley index* of $\mathrm{Inv}(N, \varphi)$ is given by

$$CH_*(N) := H_*(N, N^-)$$

the relative homology type of N and N^-. In the next section it will be shown that our computational algorithm produces such *index pairs* (N, N^-). There are a variety of results that use the Conley index to deduce information about the structure of the dynamics of the invariant set [7].

We have argued that the above approach to dynamics is well suited to numerical computations. However, a major distinction between the topological dynamics of the Conley index theory and any computational method is that the latter is necessarily combinatorial in nature. Therefore, we will now briefly repeat the above discussion but from a purely combinatorial point of view.

With this in mind, we consider the phase space to be a finite set $\mathbf{P}$, as opposed to a topological space. Furthermore, because this combinatorial system arises via an approximation, the dynamics is generated by a multivalued map $\mathbf{F}$ on this finite set. More precisely, for every $P \in \mathbf{P}$, $\mathbf{F}(P) \subset \mathbf{P}$. Note that we allow $\mathbf{F}(P) = \emptyset$. To emphasize the fact that we are considering the dynamics of multivalued maps we will write

$$\mathbf{F} : \mathbf{P} \rightrightarrows \mathbf{P}.$$

A *full trajectory through* $P \in \mathbf{P}$ is an bi-infinite sequence $\{P_n \mid n \in \mathbb{Z}\}$ satisfying $P_{n+1} \in \mathbf{F}(P_n)$ and $P_0 = P$. The *maximal invariant set* of $\mathbf{P}$ under $\mathbf{F}$ is defined to be

$$\mathrm{Inv}(\mathbf{P}, \mathbf{F}) := \{P \in \mathbf{F} \mid \exists \text{ full trajectory through } P\}$$

Let $\mathbf{A} \subset \mathbf{P}$, then

$$\mathbf{F}(\mathbf{A}) := \bigcup_{A \in \mathbf{A}} \mathbf{F}(A).$$

Inductively, $\mathbf{F}^{j+1}(P) := \mathbf{F}(\mathbf{F}^j(P))$ for every $P \in \mathbf{P}$. We will let

$$\mathbf{F}^\omega(P) := \bigcup_{j=0}^{\infty} \mathbf{F}^j(P).$$

Turning now to the dynamics of such a system. Let $P \in \mathbf{P}$ be *recurrent* if there exists $i > 0$ such that

$$P \in \mathbf{F}^i(P).$$

We can define an equivalence relation on the set of recurrent elements by

$$P \sim Q \quad \Leftrightarrow \quad \exists\, i, j > 0 \text{ such that } P \in \mathbf{F}^i(Q) \text{ and } Q \in \mathbf{F}^j(P) \tag{4}$$

As will be described in the next section, in our approximation scheme, each equivalence class (or recurrent component) will lead to an isolating neighborhood of a Morse set. To obtain an approximate Lyapunov function, we can choose any function $W : \mathbf{P} \to [0, 1]$ which satisfies the following property.

$$Q \in \mathbf{F}(P) \quad \Rightarrow \quad W(P) \geq W(Q) \quad \text{and} \quad W(P) = W(Q) \Leftrightarrow P \sim Q. \tag{5}$$

In Section 2, using a polygonal tiling of the phase space, we define a combinatorial multivalued map which approximates the flow of a differential equation. We prove that, since the multivalued map is an outer approximation of the flow, the recurrent components are isolating blocks for Morse sets of a Morse decomposition. In Section 3, we give local conditions on a triangulation which, if they can be satisfied, imply that the chain recurrent set of the flow can be approximated arbitrarily closely. In Section 4, we show some results of numerical computations.

2 Flow Transverse Polygonal Decompositions

2.1 Basic Definitions

We begin some definitions concerning simplicial complexes. While most of these definitions are standard [8], we include them for the sake of completeness.

Let $\{a_0, \ldots, a_k\}$ be a geometrically independent set in $\mathbb{R}^n$. The *k-dimensional simplex K* spanned by $a_0, \ldots, a_k$ is

$$K := \left\{ x \in \mathbb{R}^n \mid x = \sum_{i=0}^{k} t_i a_i, \text{ where } \sum_{i=0}^{n} t_i = 1 \text{ and } t_i \geq 0 \right\}$$

Any simplex spanned by a subset of $\{a_0, \ldots, a_k\}$ is called a *face* of K. The interior of K is defined by

$$\operatorname{int}(K) := \left\{ x \in \mathbb{R}^n \mid x = \sum_{i=0}^{k} t_i a_i, \text{ where } \sum_{i=0}^{k} t_i = 1 \text{ and } t_i > 0 \right\}.$$

Notice that this does not necessarily correspond to the interior of K with regard to the topology inherited from $\mathbb{R}^n$.

Definition 2.1 A *simplicial complex* $\mathcal{K}$ in $\mathbb{R}^n$ is a collection of simplicies in $\mathbb{R}^n$ satisfying:

1. Every face of a simplex of $\mathcal{K}$ is in $\mathcal{K}$.
2. The intersection of any two simplicies of $\mathcal{K}$ is a face of each of them.

Set

$$\mathcal{K}^{(l)} := \{K \in \mathcal{K} \mid \dim K = l\}$$

The *dimension* of the simplicial complex $\mathcal{K}$ is determined by its highest dimensional simplex. More precisely,

$$\dim \mathcal{K} := \max_{K \in \mathcal{K}} \dim K.$$

Observe that if $\mathcal{L}_0 \subset \mathcal{K}$ and $\mathcal{L}_1 \subset \mathcal{K}$ are simplicial complexes, then $\mathcal{L}_0 \cap \mathcal{L}_1$ is also a simplicial complex. For the purpose of this work we are only interested the following special class of simplicial complexes.

Definition 2.2 A simplicial complex $\mathcal{K}$ in $\mathbb{R}^n$ is *full* if every simplex in $\mathcal{K}$ is a face of an n-dimensional simplex in $\mathcal{K}$.

Given a full simplicial complex $\mathcal{K}$ in $\mathbb{R}^n$, $K \in \mathcal{K}^{(n-1)}$ is a *free face* if there exists a unique $L \in \mathcal{K}^{(n)}$ such that K is a face of L.

Definition 2.3 Let $\mathcal{K}$ be a full simplicial complex. Let $|\mathcal{K}|$ be the subset of $\mathbb{R}^n$ given by the union of the simplicies in $\mathcal{K}$. $|\mathcal{K}|$ is the *polytope* of $\mathcal{K}$. A *polygon* is a subset of $\mathbb{R}^n$ that is the polytope of a full simplicial complex.

Definition 2.4 Let $P = |\mathcal{P}|$ be a polygon. The *boundary* of $\mathcal{P}$ is

$$\partial\mathcal{P} := \left\{ K \in \mathcal{P}^{(n-1)} \mid K \text{ is free} \right\}$$

and the boundary of the polygon is

$$\partial P := |\partial\mathcal{P}|.$$

Definition 2.5 Let $X \subset \mathbb{R}^n$ be the polytope of the full finite simplicial complex $\mathcal{K}$. A *polygonal decomposition* of X consists of a finite collection of polygons $\{P_1, \dots P_N\}$ such that each polygon P_i is the polytope of a simplicial complex $\mathcal{P}_i \subset \mathcal{K}$,

$$X = \bigcup_{i=1}^{N} P_i,$$

and $\dim(\mathcal{P}_i \cap \mathcal{P}_j) \leq n-1$ for all $i \neq j$.

2.2 Flow Transverse Polygonal Decompositions

We are interested in the relationship between polygons and the vector field of our differential equation (1) restricted to a polygonal region X. Throughout this section $X = |\mathcal{K}|$ where $\mathcal{K}$ is a full finite simplicial complex. Our first step will be to use $\mathcal{K}$ to construct a polygonal decomposition of X which is compatible with the vector field f.

Given $K \in \mathcal{K}^{(n-1)}$, let $\nu(K)$ denote one of the two unit normal vectors to K. To determine a unique choice of $\nu(K)$, let $L \in \mathcal{K}^{(n)}$ such that K is a face of L. Then, $\nu_L(K)$ is defined to be the inward unit normal of K with respect to L.

Definition 2.6 $K \in \mathcal{K}^{(n-1)}$ is a *flow transverse face*, if

$$\nu(K) \cdot f(x) \neq 0$$

for every $x \in K$. Let $P = |\mathcal{P}|$ be a polygon where $\mathcal{P} \subset \mathcal{K}$. P is a *flow transverse polygon* if every $K \in \partial\mathcal{P}$ is flow transverse.

Let $P = |\mathcal{P}|$ be a polygon. Observe that if $K \in \partial\mathcal{P}$, then there exists a unique $K_P \in \mathcal{P}^{(n)}$ such that K is face of K_P. K is an *exit face* of P if

$$\nu_{K_P}(K) \cdot f(x) < 0 \quad \forall\, x \in K.$$

Define

$$P^- := \{K \in \partial\mathcal{P} \mid K \text{ is an exit face of } P\} \tag{6}$$

The following result follows directly from the definition of flow transversality and the fact that simplices are compact.

Lemma 2.7 *If P is a flow transverse polygon, then $f(x) \neq 0$ for all $x \in \partial P$.*

Definition 2.8 Let $\mathbf{P} = \{P_1, \ldots, P_N\}$ be a polygonal decomposition of X. Let $P_i = |\mathcal{P}_i|$. $\mathbf{P}$ is a *flow transverse polygonal decomposition* of X if the following condition is satisfied. If P_i is not flow transverse and $K \in \partial\mathcal{P}_i$ such that $\nu(K) \cdot f(x) = 0$ for some $x \in K$, then $|K| \subset \partial X$.

Observe that in the above definition, every polygon is flow transverse, except perhaps those which share an $(n-1)$-dimensional simplex with the boundary of X. Furthermore, these polygons which intersect the boundary are also flow transverse with respect to all their $(n-1)$-dimensional faces except possibly those that lie on the boundary of X. The motivation for this definition is that while we can control the structure of the polygons within X, the boundary of X is fixed and therefore for such points we have no control on the flow transversality or lack thereof.

Definition 2.9 Starting with any simplicial complex $\mathcal{K}$ and any vector field f there is a *minimal flow transverse polygonal decomposition* of X denoted by $\mathbf{P}(\mathcal{K}, f)$ which is defined as follows. Let $K_i, K_j \in \mathcal{K}^{(n)}$ such that $K_i \cap K_j = L \in \mathcal{K}^{(n-1)}$. Set $K_i \sim K_j$ if $\nu(L) \cdot f(x) = 0$ for some $x \in L$ or if $i = j$. Extend this relation by transitivity. Then $\sim$ is an equivalence relation on $\mathcal{K}$. Define $\mathbf{P}(\mathcal{K}, f) = \{P_1, \ldots, P_N\}$ to be the polygons defined by the equivalence classes. We shall let $\mathcal{P}_i \subset \mathcal{K}$ be the simplicial complex such that $P_i = |\mathcal{P}_i|$.

Our goal is to use $\mathbf{P}(\mathcal{K}, f)$ to approximate the dynamics of the flow φ restricted to X. To do this we will use the following definitions.

Definition 2.10 Two distinct polygons $P_i, P_j \in \mathbf{P}(\mathcal{K}, f)$ are *adjacent* if they share an $(n-1)$-dimensional simplex, i.e. if

$$\partial\mathcal{P}_i \cap \partial\mathcal{P}_j \cap \mathcal{K}^{(n-1)} \neq \emptyset.$$

Definition 2.11 Let $P_i, P_j \in \mathbf{P}(\mathcal{K}, f)$ be adjacent polygons and let $K \in \partial\mathcal{P}_i \cap \partial\mathcal{P}_j \subset \mathcal{K}^{(n-1)}$. *$P_i$ is in the image of P_j* if

$$\nu_{P_i}(K) \cdot f(x) > 0$$

for $x \in K$.

Definition 2.12 A polygon P is (ϵ, τ)-*recurrent* if there exists $x \in P$ such that $B_\epsilon(\varphi(\tau, x)) \cap P \neq \emptyset$. P is *non-recurrent* if it is not recurrent.

The *flow induced multivalued map* $\mathbf{F}_{\epsilon,\tau} : \mathbf{P}(\mathcal{K}, f) \rightrightarrows \mathbf{P}(\mathcal{K}, f)$ is defined as follows.

$$P_i \in \mathbf{F}_{\epsilon,\tau}(P_i) \quad \text{if and only if } P_i \text{ is } (\epsilon, \tau)\text{−recurrent}$$

and

$$P_j \in \mathbf{F}_{\epsilon,\tau}(P_i) \text{ for } j \neq i \quad \text{if and only if } P_j \text{ is in the image of } P_i.$$

In what follows, we will assume that ϵ and τ are fixed, and so to simplify the notation we will write $\mathbf{F} = \mathbf{F}_{\epsilon,\tau}$.

2.3 Local Relation between F and φ

In this section we investigate the relationship between $\mathbf{F} : \mathbf{P}(\mathcal{K}, f) \rightrightarrows \mathbf{P}(\mathcal{K}, f)$ and the flow φ generated by the differential equation (1).

As was indicated in the introduction, we will not concern ourselves with the dynamics within the polygons $P_i \in \mathbf{P}(\mathcal{K}, f)$. However, we are interested in the dynamics on the boundary points of the polygons. As above, let $P_i = |\mathcal{P}_i|$ and let

$$\mathcal{Q} = \bigcup_{i=1}^{N} \partial \mathcal{P}_i.$$

Consider $x_0 \in \operatorname{int} K$ where $K \in \partial\mathcal{P}_i \subset \mathcal{Q}^{(n-1)}$. Assume, furthermore, that $K \in \mathcal{P}_i \cap \mathcal{P}_j$. Then, $P_i \in \mathbf{F}(P_j)$ if and only if every point $x \in B_\rho(x_0) \cap K$, for ρ sufficiently small, is immediately leaving P_j and immediately entering P_i under the flow φ.

So now consider, $x_0 \in \operatorname{int} K$ where $K \in \mathcal{Q}^{(l)}$ for some $0 \le l \le n-2$. Since the polygons P_i need not be convex in general, the relationship between the dynamics of $\mathbf{F}$ and φ in a neighborhood of x_0 is more delicate. Again, consider a small ball $B_\rho(x_0)$. For the rest of this section we will assume that ρ is chosen small enough so that for any $L \in \mathcal{K}$ if $L \cap B_\rho(x_0) \neq \emptyset$ then $x_0 \in L$. For notational convenience, we will take $x_0 = 0$; in the general case the hyperplanes H discussed below should be replaced by affine spaces $x_0 + H$.

By flow transversality, K and $f(x_0)$ determine an $(l+1)$-dimensional hyperplane W. Let H^* denote the span of $W^\perp$ and $f(x_0)$, which implies $\operatorname{codim} H^* = l$. Consider any $L \in \operatorname{star}(x_0) \cap \mathcal{K}^{(n-1)}$, which necessarily contains K, and let H denote the codimension-1 hyperplane determined by L. Then H is transverse to H^* since $f(x_0) \notin H$ by flow transversality. Thus $\operatorname{codim}(H \cap H^*) = l+1$. Also, for $S_\rho = \partial B_\rho(x_0)$ define the $(n-l-1)$-dimensional sphere

$$S_\rho^* := S_\rho \cap H^*.$$

Slicing the simplicial complex with the hyperplane H^* and looking locally in $B_\rho(x_0)$, we have the following lemma.

Lemma 2.13 *Let $x_0 \in \operatorname{int}(K)$ for some $K \in \mathcal{Q}^{(l)}$ where $0 \le l \le n-2$. Then, the $(n-l-1)$-dimensional cellular complex*

$$\mathcal{S}_\rho^* := \left\{ K \cap S_\rho^* \mid K \in \operatorname{star}(x_0) \cap \mathcal{K}^{(n)} \right\}$$

is a polygonal decomposition of the sphere S_ρ^* which is in 1-1 correspondence with $\operatorname{star}(x_0) \cap \mathcal{K}^{(n)}$.

To understand the dynamics through the complex in $B_\rho(x_0)$ consider the projection of $f(x)$ onto the tangent space at x of S_ρ. This is given by

$$f_\rho(x) = f(x) - (f(x) \cdot r(x))r(x)$$

where $r(x) = (x - x_0)/\rho$. Notice that $\|r(x)\| = 1$.

Observe that for ρ small enough, the flow in $B_\rho(x_0)$ is nearly parallel. For the constant (and hence parallel) flow, $\dot{x} = f(x_0)$, the projected vector field $f(x_0) - (f(x_0) \cdot r(x))r(x)$ has exactly two zeroes which are located at the poles $x_0 \mp \rho f(x_0)/\|f(x_0)\|$ where the function

$$V(x) = -r(x) \cdot f(x_0)$$

attains its maximum and minimum values of $\pm\|f(x_0)\|$ on S_ρ. Furthermore, $V(x)$ is a Lyapunov function on S_ρ. The following lemma formalizes how these properties are preserved for f_ρ on S_ρ with ρ small.

Lemma 2.14 *Let $x_0 \in \mathcal{Q}^{(l)}$ where $0 \leq l \leq n-2$. For every $\delta > 0$, there exists $\rho_0 > 0$ such that for $0 < \rho < \rho_0$:*

1. *if $x \in S_\rho$ is a zero of f_ρ, then $\|f(x_0)\| - |V(x)| < \delta$, and*
2. *if $x \in S_\rho \setminus \{x : \|f(x_0)\| - |V(x)| < \max\{\delta, 2\delta/\|f(x_0)\|\}\}$, then*
$$\dot{V}(x) = -f_\rho(x) \cdot f(x_0) \leq -\delta < 0.$$

Proof: First choose $\alpha > 0$ so that $\|f(x_0)\| - |y| < \delta$ whenever $\|f(x_0)\|^2 - y^2 < \alpha$ and $|y| \leq \|f(x_0)\|$. By continuity, ρ_0 can be chosen so that

$$\|f(x) - f(x_0)\| < \min\{\alpha/2\|f(x_0)\|, \delta\}$$

and

$$|f(x) \cdot r(x) - f(x_0) \cdot r(x)| \leq \|f(x) - f(x_0)\| \, \|r(x)\| < \min\{\alpha/2\|f(x_0)\|, \delta\}$$

for all $\rho < \rho_0$. If x is a zero of f_ρ, then $f(x) = (f(x) \cdot r(x))r(x)$ which implies

$$\begin{aligned} \|(f(x_0) \cdot r(x))r(x) - f(x_0)\| &\leq \|(f(x_0) \cdot r(x))r(x) - (f(x) \cdot r(x))r(x)\| \\ &\quad + \|f(x) - f(x_0)\| < \alpha/\|f(x_0)\|. \end{aligned}$$

Taking the inner product with $f(x_0)$ gives $\|f(x_0)\|^2 - (r(x) \cdot f(x_0))^2 < \alpha$ which by the choice of α implies (1).

Moreover,

$$-f_\rho(x) \cdot f(x_0) = -f(x) \cdot f(x_0) + (f(x) \cdot r(x))(f(x_0) \cdot r(x)).$$

Again $\|f(x) - f(x_0)\| < \delta$ implies $-f(x) \cdot f(x_0) < -\|f(x_0)\|^2 + \delta$, and $|f(x) \cdot r(x) - f(x_0) \cdot r(x)| < \delta$ implies $|f(x) \cdot r(x)| < |f(x_0) \cdot r(x)| + \delta$. Hence

$$\begin{aligned} |f(x) \cdot r(x)| \cdot |f(x_0) \cdot r(x)| &< (|f(x_0) \cdot r(x)| + \delta)|f(x_0) \cdot r(x)| \\ &< (\|f(x_0)\| - \delta + \delta)(\|f(x_0)\| - \delta_*) \\ &< \|f(x_0)\|^2 - 2\delta \end{aligned}$$

where the second inequality follows from

$$x \in S_\rho \setminus \{x : \|f(x_0)\| - |V(x)| < \max\{\delta, 2\delta/\|f(x_0)\|\}\}.$$

Therefore $-f_\rho(x) \cdot f(x_0) < -\|f(x_0)\|^2 + \delta + \|f(x_0)\|^2 - 2\delta = -\delta$, which proves (2). ∎

The last statement says that V is a Lyapunov function for the flow $\dot{x} = f_\rho(x)$ on $S_\rho \setminus \{x : \|f(x_0)\| - |V(x)| < \max\{\delta, 2\delta/\|f(x_0)\|\}\}$.

Lemma 2.15 *Let $x_0 \in \operatorname{int}(K)$ for some $K \in \mathcal{Q}^{(l)}$ where $0 \leq l \leq n-2$. There exists a $\rho_1 > 0$ such that for all $0 < \rho < \rho_1$ we have*

1. *$\mathcal{S}^*_\rho$ is transverse to the flow f^*_ρ on S^*_ρ,*
2. *the flow induced multivalued map $\mathbf{F}^*_\rho : \mathcal{S}^*_\rho \rightrightarrows \mathcal{S}^*_\rho$ is locally equivalent to the flow induced map $\mathbf{F}$.*

Proof: Let $f^*(x) = P_{H^*} f(x)$ be the projection of $f(x)$ onto the hyperplane H^*. Since $f(x_0) \in H^*$, we can apply Lemma 2.14 to the vector field f^* on H^* and consider the vector field f^*_ρ on S^*_ρ. Any $(n-l-2)$-dimensional face L^* in $\mathcal{S}^*_\rho$ is the intersection of S^*_ρ with a hyperplane $H \cap H^*$ where H is determined by a cell in $\mathcal{K}^{(n-1)} \cap \operatorname{star}(x_0)$ as described above.

For any point $x \in L^*$, the vector field $f^*(x) \notin H$ by flow transversality, and hence $f^*_\rho(x) = f^*(x) - (f^*(x) \cdot r(x))r(x) \notin H$ since $r(x) \in H$. ∎

Definition 2.16 Let $x_0 \in \operatorname{int} K$ where $K \in \partial \mathcal{P}_i$. Let

$$\mathcal{L} := \left\{ L \in \mathcal{P}_i^{(n)} \mid K \text{ is a face of } L \right\}.$$

P_i is *field enclosing for f at x_0* if

$$f(x_0) \cdot \nu(L) > 0$$

for every $L \in \mathcal{L}$ where $\nu(L)$ is the inward unit normal vector to L with respect to P_i.

Lemma 2.17 *Let $x_0 \in \operatorname{int} K \cap \operatorname{int} X$ where $K \in \mathcal{Q}^{(l)}$ for some $0 \leq l \leq n-2$. Then there exist unique elements A and R in $\mathbf{P}(\mathcal{K}, f)$ such that A is field-enclosing for f at x_0 and R is field enclosing for $-f$ at x_0. Moreover, for every $P_i \in \mathbf{P}(\mathcal{K}, f)$ such that $P_i \cap B_\rho(x_0) \neq \emptyset$, $A \in \mathbf{F}^\omega(P_i)$ and $P_i \in \mathbf{F}^\omega(R)$.*

Proof: The existence of the polyhedra A and R with the stated enclosure properties follows immediately from flow transversality. The nontrivial assertion of the lemma is that $A \in \mathbf{F}^\omega(P_i)$ for every $P_i \in \mathbf{P}(\mathcal{K}, f)$ such that $P_i \cap B_\rho(x_0) \neq \emptyset$.

Then applying that result to the backward flow, given by the vector field $-f$, and noting that $P \in F^\omega(Q)$ iff $Q \in (-F)^\omega(P)$, we obtain $P \in F^\omega(R)$ for every $P \cap B_\rho(x) \neq \emptyset$.

First, we choose $\rho > 0$ small enough so that Lemma 2.15 applies. Let $R^* = R \cap S^*_\rho$ and $A^* = A \cap S^*_\rho$. For any $\rho > 0$, R^* and A^* contain sectors around $\pm f(x_0)$ of fixed angular size, i.e. by choosing δ small enough, $\{x \in S^*_\rho : \|f^*(x_0)\| - |V(x)| > \delta\} \subset \operatorname{int}(R^* \cup A^*)$. Then, by Lemma 2.14, we have that $V(x)$ is a Lyapunov function on $S^*_\rho \backslash \operatorname{int}(R^* \cup A^*)$.

Step 1: If $\sigma \in \mathcal{S}^*_\rho \backslash \{A^*\}$, then $\min_{x \in \sigma} V(x)$ is attained at a vertex of σ.

The set $H_{\min} = \{x \in H^* : V(x) = \min_{x \in \sigma} V(x)\}$ is a $(n-l-1)$-dimensional affine space of the form $x_{\min} + \operatorname{span}\{f(x_0)\}^\perp$. Furthermore, $H_{\min}$ is tangent to the sphere S^*_ρ only at the global minimum of V which is attained at the pole in A^*, and thus for $\sigma \neq A^*$, we have that $H_{\min}$ intersects S^*_ρ transversely.

If $H_{\min} \cap \operatorname{int}(\sigma) \neq \emptyset$ then σ must contain points for which V is less than $\min_{x \in \sigma} V(x)$, and hence $H_{\min} \cap \sigma \subset \partial\sigma$. The set $\partial\sigma$ is composed of the intersection of S^*_ρ with $(n-l-1)$-dimensional hyperplanes. Since $H_{\min} \cap \operatorname{int}(\sigma) \neq \emptyset$, $H_{\min}$ must contain a complete face of σ and hence a vertex.

Step 2: If $\sigma \in \mathcal{S}_\rho^* \backslash \{A^*\}$, then there exists $\psi \in (\mathbf{F}_\rho^*)^\omega(\sigma)$ such that $\min_{x\in\psi} V(x) < \min_{x\in\sigma} V(x)$.

Let $v \in \mathcal{S}_\rho^*$ be a vertex at which $V(x)$ attains its minimum in σ. Since $V(x)$ decreases in the direction $f_\rho^*(v)$, the element σ cannot be field-enclosing at v. By flow transversality, there exists $\psi_1 \in \mathbf{F}_\rho^*(\sigma)$ such that $v \in \psi_1$. If $\min_{x\in\psi_1} V(x) < \min_{x\in\sigma} V(x)$, then we have proven the claim, otherwise $\min_{x\in\psi_1} V(x) = \min_{x\in\sigma} V(x)$, which is again attained at v. In this case, ψ_1 cannot be field-enclosing at v, and we can repeat the previous step to obtain $\psi_2 \in \mathbf{F}_\rho^*(\psi_1)$, which implies $\psi_2 \in (\mathbf{F}_\rho^*)^2(\sigma)$. Since there are only finitely many elements in $\mathcal{S}_\rho^*$, this process must terminate and yield $\psi \in (\mathbf{F}_\rho^*)^\omega(\sigma)$ such that $\min_{x\in\psi} V(x) < \min_{x\in\sigma} V(x)$.

Step 3: If $\sigma \in \mathcal{S}_\rho^* \backslash \{A^*\}$, then $A^* \in (\mathbf{F}_\rho^*)^\omega(\sigma)$.

If $\sigma = A^*$, then $A^* \in F^\omega(A^*)$, and hence we assume that $\sigma \neq A^*$. From Step 2 we can find $\psi_1 \in (\mathbf{F}_\rho^*)^\omega(\sigma)$ such that $\min_{x\in\psi_1} V(x) < \min_{x\in\sigma} V(x)$. If $\psi_1 \neq A^*$, then we can repeat the previous step to obtain $\psi_2 \in \mathbf{F}_\rho^*(\psi_1)$, which implies $\psi_2 \in (\mathbf{F}_\rho^*)^2(\sigma)$, and $\min_{x\in\psi_2} V(x) < \min_{x\in\psi_1} V(x)$. Since there are only finitely many elements in $\mathcal{S}_\rho^*$, this process must terminate, at which point $A^* = \psi \in (\mathbf{F}_\rho^*)^\omega(\sigma)$.

Step 3 immediately implies the desired result by the correspondences between the maps $\mathbf{F}$ and $\mathbf{F}_\rho^*$ and the sets $\mathrm{star}(x_0) \cap \mathcal{K}^{(n)}$ and $\mathcal{S}_\rho^*$ in Lemma 2.15. ∎

2.4 Global Relation between F and φ

Recall that we are studying the dynamics of (1) restricted to the polygon X where $X = |\mathcal{K}|$. As was indicated in Section 2.2 this gives rise to the polygonal decomposition $\mathbf{P}(\mathcal{K}, f)$ and the induced multivalued map $\mathbf{F} : \mathbf{P}(\mathcal{K}, f) \rightrightarrows \mathbf{P}(\mathcal{K}, f)$. We are now in a position to compare the global dynamics of φ with that of $\mathbf{F}$.

First we observe that, given a typical point $x \in X$, there is no reason to believe that its entire forward trajectory lies in X. Therefore, define

$$\tau_x := \max\{t \geq 0 \mid \varphi([0,t],x) \subset X\}$$

The first theorem indicates that forward trajectories of $\mathbf{F}$ provide an outerbound for trajectories of φ.

Theorem 2.18 *Let $x_0 \in X$. Let $P \in \mathbf{P}(\mathcal{K}, f)$ such that $x_0 \in P$. Then*

$$\varphi(x_0, (0, \tau_{x_0})) \subset \mathrm{int}(|\mathbf{F}^\omega(P)|)$$

Proof: We will show that for any $t \in [0, \tau_{x_0}]$, if $\varphi(x_0, t) \in |\mathbf{F}^\omega(P)| \cap \mathrm{int}(X)$, then there is an $\epsilon > 0$ such that $t + \epsilon \in [0, \tau_{x_0}]$ and $\varphi(x_0, [t, t+\epsilon]) \subset |\mathbf{F}^\omega(P)|$. The result then follows by considering $\tau = \inf\{t \in [0, \tau_{x_0}] \mid \varphi(x_0, t) \notin |\mathbf{F}^\omega(P)|\}$. If τ exists, then $\varphi(x_0, \tau) \in \partial X$.

To prove the above claim, we consider three cases. First, if $\varphi(x_0, t) \in \mathrm{int}(P')$ for some $P' \in \mathbf{F}^\omega(P)$, then the desired ϵ exists since $\mathrm{int}(P')$ is open. Second, if $\varphi(x_0, t) \in \mathrm{int}(K)$ for some $\sigma \in \mathcal{Q}^{(n-1)}$, then by flow transversality, there exist elements P_i and P_j in $\mathbf{P}(\mathcal{K}, f)$ such that $K \in \partial\mathcal{P}_i \cap \partial\mathcal{P}_j$ and $P_i \in \mathbf{F}(P_j)$. By definition, there exists $\epsilon > 0$ such that

$\varphi(x_0, [t, t+\epsilon]) \subset P_i$. We are assuming that $\varphi(x_0, t) \in |\mathbf{F}^\omega(P)|$ so that either $P_i \in \mathbf{F}^\omega(P)$ or $P_j \in \mathbf{F}^\omega(P)$, which implies $P_i \in \mathbf{F}^\omega(P)$.

The remaining case is that $\varphi(x_0, t) \in \operatorname{int}(K)$ for some $K \in \mathcal{Q}^l$ with $0 \leq l \leq n-2$. We are assuming that $\varphi(x_0, t) \in |\mathbf{F}^\omega(P)|$, which implies that $\varphi(x_0, t) \in P'$ for some P' such that $K \subset P'$ and $K \in \operatorname{star}(\varphi(x_0, t)) \cap \mathbf{F}^\omega(P)$. Let $A \in \mathbf{P}(\mathcal{K}, f)$ be the element which is field-enclosing at $\varphi(x_0, t)$. Then, there exists $\epsilon > 0$ such that $\varphi(x_0, [t, t+\epsilon]) \subset A$. By Lemma 2.17, we have $A \in \mathbf{F}^\omega(P') \subset \mathbf{F}^\omega(P)$, which completes the proof. ∎

We need to worry about internal tangencies. To eliminate these we introduce the following definition.

Definition 2.19 Let $x_0 \in \partial X$ such that $X \cap B_\rho(x_0)$ is not convex for arbitrarily small $\rho > 0$. Then, $x_0 \in K_1 \cap K_2$ where $K_i \in \partial \mathcal{P}$ are distinct $(n-1)$-dimensional simplexes. Assume that K_i is a face of $L_i \in \mathcal{P}_i$ (it is possible that $P_1 = P_2$). Then x_0 is a *re-entry point* if

$$\nu_{L_1}(K_1) \cdot f(x_0) > 0 \quad \text{and} \quad \nu_{L_2}(K_2) \cdot f(x_0) < 0.$$

A collection of polygons $\mathbf{Q} \subset \mathbf{P}(\mathcal{K}, f)$ is *re-entry free* if

1. every $(n-1)$-dimensional simplex in $\partial X \cap \mathbf{Q}$ is flow transverse, and
2. $\mathbf{Q}$ contains no re-entry points.

Theorem 2.20 *Let $\mathbf{Q} \subset \mathbf{P}(\mathcal{K}, f)$ be an equivalence class (4) of the recurrent set of $\mathbf{F}$. If $\mathbf{Q}$ is re-entry free, then $(\mathbf{Q}, \mathbf{Q}^-)$ is an index pair.*

Proof: To prove that $(\mathbf{Q}, \mathbf{Q}^-)$ is an index pair it needs to be shown that condition (3) is satisfied for every $x \in \partial \mathbf{Q}$. The proof is by contradiction. So assume there exists $x_0 \in \partial \mathbf{Q}$ and an $\epsilon > 0$ such that

$$\varphi((-\epsilon, \epsilon), x_0) \subset \mathbf{Q}. \tag{7}$$

The assumption that $\mathbf{Q}$ is re-entry free implies that $x_0 \notin \partial X$. By construction the $(n-1)$-faces of each Q_i are flow transverse, and therefore we can further assume that $x_0 \in |\mathcal{Q}^{(j)}|$ for some $0 \leq j \leq n-2$.

Combining (7) and Lemma 2.17 we can conclude that there exist unique (though not necessarily distinct) polygons A and R such that if $P_i \in \mathbf{P}(\mathcal{K}, f)$ and $x_0 \in P_i$, then $A \in \mathbf{F}^\omega(P_i)$ and $P_i \in \mathbf{F}^\omega(R)$. However, A and R are related by (4). Thus every P_i containing x_0 is related to A and R by (4), and hence every such $P_i \in \mathbf{Q}$. This contradicts the assumption that $x_0 \in \partial \mathbf{Q}$. ∎

A similar proof leads to the next corollary. For the remainder of this section let

$$\mathbf{S} := \operatorname{Inv}(\mathbf{P}(\mathcal{K}, f), \mathbf{F})$$

Corollary 2.21 *If $\mathbf{S}$ is re-entry free, then $\operatorname{Inv}(X, \varphi)$ is an isolated invariant set.*

Observe that it follows from the proof of Theorem 2.20 that any equivalence class in the recurrent set of $\mathbf{F}$ which does not touch the boundary of X is automatically re-entry free.

Let $\{\mathbf{Q}(j) \mid j \in \mathcal{J}\}$ be the set of equivalence classes of the recurrent set of $\mathbf{F}$ and let $W : \mathbf{P}(\mathcal{K}, f) \to [0, 1]$ be an approximate Lyapunov function satisfying (5). Let $>$ be a partial ordering of $\mathcal{J}$ defined by

$$W(\mathbf{Q}(j)) > W(\mathbf{Q}(k)) \quad \Rightarrow \quad j > k$$

Theorem 2.22 *Let* $M(j) := \mathrm{Inv}(\mathbf{Q}(j), \varphi)$*. Then*

$$\{M(j) \mid j \in \mathcal{J}\}$$

is a Morse decomposition of $\mathrm{Inv}(X, \varphi)$ *with admissible order* $>$*.*

The following lemma will be used in the proof of this theorem.

Lemma 2.23 *If* $x \in \mathrm{Inv}(X, \varphi)$*, then* $\omega(x) \subset M(q)$ *and* $\alpha(x) \subset M(p)$ *for some* $q \in \mathcal{J}$*.*

Proof: We shall prove the first case since the second is similar. The proof is by contradiction. Assume that there exists $P \in \mathbf{P}(\mathcal{K}, f)$ such that $\omega(x) \cap P \neq \emptyset$ and $P \notin \mathbf{Q}(i)$ for all $i \in \mathcal{J}$. This latter condition implies that $P \notin \mathbf{F}^n(P)$ for all $n \geq 1$. Let $y \in P \cap \omega(x)$ and let $\{t_m\}$ be an increasing sequence of positive times such that $\lim_{m\to\infty} \varphi(t_m, x) = y$.

By Theorem 2.18 and the fact that $P \notin \mathbf{F}^n(P)$ for all $n \geq 1$ we can assume that $\varphi(t_m, x) \notin P$ for all m. This implies that $y \in \partial P$. Assume that there exists an $\epsilon > 0$ for which $\varphi((-\epsilon, 0], y) \subset P$. Since P is flow transverse $\varphi((-\epsilon, 0), y) \subset \mathrm{int}(P)$. This would imply that we could choose a sequence of t_m such that $\varphi(t_m, x) \in P$ a contradiction. A similar argument applies to the case that $\varphi([0, \epsilon), y) \subset P$. Therefore, we can conclude that there exists $\epsilon > 0$ such that $\varphi((-\epsilon, \epsilon), y) \cap P = y$. This in turn implies that $y \in \mathcal{Q}^{(k)}$ for some $0 \leq k \leq n-2$.

We can now conclude the existence of $A, R \in \mathbf{P}(\mathcal{K}, f)$ as in Lemma 2.17. Since

$$\varphi((-\epsilon, 0], y) \subset R \quad \text{and} \quad \varphi([0, \epsilon), y) \subset A,$$

the previous arguments imply that A and R belong to some $\mathbf{Q}(j)$. But the fact that $A \in F^\omega(P)$ and $P \in F^\omega(R)$ implies that $P \in \mathbf{Q}(j)$, the desired contradiction. ∎

Proof of Theorem 2.22: Let $x \in \mathrm{Inv}(X, \varphi)$. If $\varphi(\mathbb{R}, x) \subset \mathbf{Q}(j)$ for some $j \in \mathcal{J}$, then $x \in M(j)$. Now consider $x \in \mathrm{Inv}(X, \varphi) \setminus \cup M(j)$. By Lemma 2.23, $\omega(x) \subset M(q)$ and $\alpha(x) \subset M(p)$ for some $p, q \in \mathcal{J}$. We need to show that $p > q$. By definition, there exists $P \in \mathbf{P}(\mathcal{K}, f)$ and s such that $P \notin \mathbf{Q}(q)$ and $\varphi(s, x) \in P$. By Theorem 2.18, $\mathbf{Q}(q) \subset \mathbf{F}^\omega(P)$. However, $P \notin \mathbf{F}^\omega(\mathbf{Q}(q))$. Thus, $W(P) > W(R)$ for any $R \in \mathbf{Q}(q)$. Clearly, if $Q \in \mathbf{Q}(p)$ then $W(Q) \geq W(P)$. Thus, $p > q$. ∎

Given that we have a Morse decomposition, the next step is to produce an index filtration. From an index filtration the Conley indices of all Morse sets associated with the Morse decomposition can be computed. However, before we can proceed several definitions are needed.

Definition 2.24 Given a Morse decomposition $\{M(j) \mid j \in (\mathcal{J}, >)\}$ an *index filtration* is a collection of compact sets

$$\mathcal{N} := \{N(I) \mid I \in \mathcal{A}(\mathcal{J}, >)\}$$

satisfying the following two conditions:

1. for each $I \in \mathcal{A}(\mathcal{J}, >)$, $(N(I), N(\emptyset))$ is an index pair for $M(I)$.

2. for each $I, J \in \mathcal{A}(\mathcal{J}, >)$,

$$N(I \cap J) = N(I) \cap N(J) \quad \text{and} \quad N(I \cup J) = N(I) \cup N(J).$$

Returning to the polygonal setting of this section define

$$N(\emptyset) := \mathbf{S}^-.$$

For $I \in \mathcal{A}(\mathcal{J}, >)$ define

$$N(I) := \left(\mathbf{S} \cap \bigcup_{j \in I} \mathbf{F}^\omega(\mathbf{Q}(j))\right) \cup N(\emptyset).$$

Theorem 2.25 *If* S *is re-entry free, then*

$$\mathcal{N} := \{N(I) \mid I \in \mathcal{A}(\mathcal{J}, >)\}$$

is an index filtration for the Morse decomposition $\{M(j) \mid j \in \mathcal{J}\}$

Proof: The second condition of Definition 2.24 is obviously satisfied, therefore we only need consider the first.

Given $x \in \partial N(I)$ we need to show that for every $\epsilon > 0$, $\varphi((-\epsilon, \epsilon), x) \not\subset N(I)$. The proof is by contradiction. So assume $x_0 \in \partial N(I)$ and $\varphi((-\epsilon, \epsilon), x_0) \subset N(I)$. Since $\mathbf{S}$ is re-entry free, $x_0 \notin \partial X$.

By Lemma 2.17 there exists $A, R \in \mathbf{P}(\mathcal{K}, f)$ with the property that for every $P_i \in \mathbf{P}(\mathcal{K}, f)$ such that $x_0 \in P_i$, $P_i \in \mathbf{F}^\omega(R)$ and $\varphi((-\epsilon, \epsilon), x_0) \subset A \cup R$. The last condition implies that $R \in N(I)$. By definition, $R \in N(I)$ implies that $\mathbf{F}^\omega(R) \subset N(I)$. Therefore, if $x_0 \in P_i \in \mathbf{P}(\mathcal{K}, f)$ then $P_i \in N(I)$. However, this contradicts the assumption that $x_0 \in \partial N(I)$. ∎

3 Polygonal Approximation of Chain Recurrence

In the previous section we showed that given any triangulation one can recover a Morse decomposition and an associated index filtration. However, no assumptions were made on the triangulation, and hence the approximation can be arbitrarily bad. In fact, in the worst case it is possible that $|\mathbf{P}(\mathcal{K}, f)| = X$. In this section we will give a local condition on the triangulation which guarantees that the global dynamics of the multivalued map on polygons is a good approximation to the global dynamics of the flow. Our goal is to show that for any $\epsilon > 0$ there is a numerically computable, local orientation condition which, when imposed on each simplex, implies that the resulting Morse decomposition is at least as fine as the decomposition by ϵ-chain recurrent components given in Theorem 1.4.

In this section, we will assume that $\mathcal{K}$ is a full, finite simplicial complex for which $|\mathcal{K}| = X$ contains no equilibrium points of the flow. In this situation, the trajectories of $x' = f(x)/\|f(x)\|$ are reparametrizations of the trajectories of $x' = f(x)$ by arclength, which does not affect the global dynamics. Hence without loss of generality, the vector field f will be assumed to be a unit vector field.

3.1 Local Orientation

For convenience, we will use the terminology σ is a *facet* of K to describe the situation that $K \in \mathcal{K}^{(n)}$ is a highest dimensional simplex and σ is an $(n-1)$-dimensional face in ∂K. If σ is a facet of K, then there is one vertex v of K which is not a vertex of σ, and we will find it convenient to denote $K = [\sigma, v]$. It will be important to study the direction of the vector field on the facets of simplices which leads to the following definitions.

Definition 3.1 A facet σ of K is an *infacet* if there exists $x \in \sigma$ such that $f(x) \cdot \nu_K(\sigma) \geq 0$.

Definition 3.2 For $\delta > 0$, a simplex $K \in \mathcal{K}^{(n)}$ is *δ-oriented* with respect to the unit vector field f if each infacet σ of K satisfies the property that if $K = [\sigma, v]$, then there exists $y \in \sigma$ such that $(v-y) \cdot f(y) \geq \delta \|v - y\|$. A simplicial complex $\mathcal{K}$ is *δ-oriented* if each $K \in \mathcal{K}^{(n)}$ is δ-oriented.

Many of the geometric estimates required in this section, including the above definition, can be given in terms of cones, which will be denoted as follows. For $v \in S^{n-1}$ the (positive) cone with axis along v and angle $\cos^{-1}(\alpha)$ is $C(v,\alpha) = \{u \in \mathbb{R}^n : u \cdot v \geq \alpha \|u\|\}$, and for $x_0 \in \mathbb{R}^n$ the affine cone is given by $C(x_0, v, \alpha) = x_0 + C(v, \alpha)$. Thus a simplex is δ-oriented if for every infacet σ with opposite vertex v, there exists $y \in \sigma$ such that $v - y \in C(f(y), \delta)$.

We begin with two technical lemmas about cones.

Lemma 3.3 *Let $f, g, h \in S^{n-1}$. Let $1 \geq \omega > \lambda > 0$ and $1 \geq \delta \geq \Delta(\omega, \lambda) := \omega\lambda + \sqrt{(1-\lambda^2)(1-\omega^2)}$. If $f \in C(g, \omega)$ and $h \in C(f, \delta)$, then $h \in C(g, \lambda)$.*

Proof: The cone assumptions amount to $f \cdot g \geq \omega$ and $h \cdot f \geq \delta$ then $h \cdot g \geq \lambda$. Without loss of generality we can assume $f = e_1$ by applying a rotation. Then we have that

$$\begin{array}{ccccccc} f \cdot g & = & g_1 & = & \cos(\theta_1) & \geq & \omega \\ h \cdot f & = & h_1 & = & \cos(\theta_3) & \geq & \delta \end{array}$$

and let $h \cdot g = \cos(\theta_2)$. It follows that

$$\begin{array}{ccccccc} \sum_{i=2}^n g_i^2 & = & 1 - g_1^2 & = & \sin^2(\theta_1) & \leq & 1-\omega^2 \\ \sum_{i=2}^n h_i^2 & = & 1 - h_1^2 & = & \sin^2(\theta_3) & \leq & 1-\delta^2 \\ \sum_{i=1}^n h_i g_i & = & h_1 g_1 + \sum_{i=2}^n h_i g_i & = & \cos(\theta_2) & \geq & h_1 g_1 - |\sum_{i=2}^n h_i g_i|. \end{array}$$

By Cauchy-Schwartz we get

$$\begin{aligned} \cos(\theta_2) \geq h_1 g_1 - (\sum_{i=2}^n h_i^2)^{1/2} (\sum_{i=2}^n g_i^2)^{1/2} &= \cos(\theta_1)\cos(\theta_3) - \sin(\theta_1)\sin(\theta_3) \\ &\geq \omega\delta - \sqrt{1-\omega^2}\sqrt{1-\delta^2} \end{aligned}$$

Now if we impose the condition

$$\omega\delta - \sqrt{1-\omega^2}\sqrt{1-\delta^2} \geq \lambda, \tag{8}$$

then $\cos(\theta_2) \geq \lambda$ as we wish. Solving the inequality (8) for δ leads to the inequality

$$1 \geq \delta \geq \omega\lambda + \sqrt{(1-\lambda^2)(1-\omega^2)}.$$

■

Lemma 3.4 *Let* $0 < \Lambda < \lambda < 1$ *and* $g \in S^{n-1}$*. For any unit vectors* $\mu \in C(g, \Lambda)$ *and* $\nu \in C(g, \lambda)$*,*

$$\|\mu - \nu\| \leq \sqrt{2}(1 - \Lambda\lambda + \sqrt{1 - \Lambda^2}\sqrt{1 - \lambda^2})^{1/2}. \tag{9}$$

Proof: Let $\theta_1 = \cos^{-1}(\Lambda)$ and $\theta_2 = \cos^{-1}(\lambda)$. Then the largest angle between any vector in the cone $C(g, \Lambda)$ and any vector in the cone $C(g, \lambda)$ is $\theta_1 + \theta_2$. Therefore $\mu \cdot \nu \geq \cos(\theta_1 + \theta_2) = \Lambda\lambda - \sqrt{1 - \Lambda^2}\sqrt{1 - \lambda^2}$. Thus $\|\mu - \nu\|^2 = 2 - 2\mu \cdot \nu \leq 2(1 - \Lambda\lambda - \sqrt{1 - \Lambda^2}\sqrt{1 - \lambda^2})$. ■

The next lemma implies that on a complex oriented to a vector field which lies in a cone, the multivalued map also respects cones, i.e. if a simplex lies in a suitable cone, then so must the entire polygon which contains it and the entire forward image of that polygon.

Lemma 3.5 *Let* $\omega, \lambda, \delta > 0$ *satisfy* $1 \geq \delta \geq \Delta(\omega, \lambda)$ *and* $1 \geq \omega > \lambda$*. Suppose* $\mathcal{K}$ *is a* δ*-oriented complex with respect to the vector field* f *which maps into a cone* $C(g, \omega)$ *for some* $g \in S^{n-1}$*. Suppose* $K, L \in \mathcal{K}^{(n)}$ *and* $K \cap L = \sigma$ *is an infacet of* L*. If* $K \subset C(p, g, \lambda)$ *for some* $p \in \mathbb{R}^n$ *then* $L \subset C(p, g, \lambda)$*.*

Proof: Let $\zeta = C(p, g, \lambda)$, $L = [\sigma, v]$ and $\nu = \nu_L(\sigma)$. Since σ is a facet of K, it lies entirely in the cone ζ. Moreover, since L is δ-oriented, there exists $y \in \sigma$ such that

$$(v - y) \cdot f(y) \geq \delta\|v - y\|,$$

i.e. $(v - y) \in C(f(y), \delta)$. We have also assumed that $f(y) \in C(g, \omega)$. These relations together with the conditions on ω, δ, and λ imply by Lemma 3.3 that $v - y \in C(g, \lambda)$.

Since $y \in K \subset \zeta$ and ζ is a cone, the half line given by

$$y + t\frac{(v - y)}{\|v - y\|}$$

is contained in ζ for all $t \geq 0$. Choosing $t = \|v - y\|$ implies $v \in \zeta$. Since L is the convex hull of σ and v, both of which are contained in ζ, we have $L \subset \zeta$. ■

Corollary 3.6 *Let* $\omega, \lambda, \delta > 0$ *satisfy* $1 \geq \delta \geq \Delta(\omega, \lambda)$ *and* $1 \geq \omega > \lambda$*. Suppose* $\mathcal{K}$ *is a* δ*-oriented complex with respect to the vector field* f *which maps into a cone* $C(g, \omega)$ *for some* $g \in S^{n-1}$*. Let* $K \in \mathcal{K}^{(n)}$ *with* $K \subset P \in \mathbf{P}(\mathcal{K}, f)$*. If* $K \subset C(p, g, \lambda)$ *for some* $p \in \mathbb{R}^n$ *then* $P \subset C(p, g, \lambda)$ *and* $|\mathbf{F}^\omega(P)| \subset C(p, g, \lambda)$*.*

Proof: By definition of $\mathbf{P}(\mathcal{K}, f)$, if $L \in \mathcal{K}^{(n)}$ and $L \subset P$, then there exist simplices $K_0, \ldots, K_r$ such that $K = K_0$, $L = K_r$, and for $i = 1, \ldots, r$ we have $K_{i-1} \cap K_i = \sigma_i$ is a facet of K_i with $f(x) \cdot \nu(K_i) = 0$ for some $x \in \sigma_i$, which implies that σ_i is an infacet of K_i. Repeated application of Lemma 3.5 yields $K_i \subset \zeta$ for all $i = 0, \ldots, r$ and hence $L \subset \zeta$. Therefore, $P \subset \zeta$ since L is an arbitrary simplex in P.

Similarly, if $Q \in F^\omega(P)$ and $L \in \mathcal{K}^{(n)}$ with $L \subset Q$, then by the definition of the multivalued map $\mathbf{F}$, there are simplices $K_0, \ldots, K_r$ such that $K_0 \subset P$, $K_r = L$, and $K_{i-1} \cap K_i = \sigma_i$ is an infacet of K_i for $i = 1, \ldots, r$. Lemma 3.5 then implies $Q \subset \zeta$. ■

Lemma 3.7 *Let $\omega, \lambda, \delta > 0$ satisfy $1 \geq \delta \geq \Delta(\omega, \lambda)$ and $1 \geq \omega > \lambda$. Suppose $\mathcal{K}$ is a δ-oriented complex with respect to the vector field f which maps into a cone $C(g, \omega)$ for some $g \in S^{n-1}$. Let $\{K_i\}_{i=-}^{m}$ be simplices for which $K_0 \subset C(p, g, \lambda)$ and $K_{i-1} \cap K_i = \sigma_i$ is an infacet of K_i and v_i be a vertex of K_i satisfying $v_i \cdot g = \min_{v \in K_i \cap \mathcal{K}^{(0)}} v \cdot g$. Then $v_i \cdot g$ is nondecreasing.*

Proof: By definition of v_i we have $v_i \cdot g \leq y \cdot g$ for all $y \in K$. Since $\mathcal{K}$ is δ-oriented, there exists $y \in \sigma_i$ such that $(v_{i+1} - y) \cdot f(y) \geq \delta \|v_{i+1} - y\|$. Therefore

$$(v_{i+1} - v_i) \cdot g \geq (v_{i+1} - y) \cdot g \geq \lambda \|v_{i+1} - y\| > 0$$

by Lemma 3.3 since $f \in C(g, \omega)$. ∎

The next lemma shows that any oriented simplex is contained in some cone whose apex is a computable distance away. This distance depends on the orientation of the cone and the projection of the simplex orthogonal to the cone axis. Let $g \in S^{n-1}$ and $K \in \mathcal{K}^{(n)}$. Define $H_g(v) = v + \{x \in \mathbb{R}^n : x \cdot g = 0\}$. Then there exists a vertex v of K such that

$$K \subset H_g^+(v) = v + \{x \in \mathbb{R}^n : x \cdot g \geq 0\}.$$

Let $H_g = H_g(v)$ and $H_g^+ = H_g^+(v)$ for this choice of v. Thus, H_g is a supporting hyperplane of K such that K is contained in the half-space H_g^+. Let $\pi_g : \mathbb{R}^n \to H_g$ be the orthogonal projection onto H_g. For any simplex σ, denote by $\rho_g(\sigma)$ the diameter of the projection $\pi_g \sigma$, and define $\rho_g(K) := \max\{\rho_g(\sigma) : \sigma \in \partial K \text{ is an infacet of } K\}$ for $K \in \mathcal{K}^{(n)}$.

Lemma 3.8 *Let ω, δ, λ satisfy $1 \geq \delta \geq \Delta(\omega, \lambda)$ and $1 \geq \omega > \lambda > 0$. Suppose $\mathcal{K}$ is a δ-oriented complex with respect to the vector field f which maps into a cone $C(g, \omega)$ for some $g \in S^{n-1}$. Then there exists a point $p \in \mathbb{R}^n$ such that $\operatorname{dist}(p, H_g) = \rho_g(K)\lambda/\sqrt{1 - \lambda^2}$ and $K \subset C(p, g, \lambda)$.*

Proof: Let v be the vertex of K in H_g. Since $f \in C(g, \omega)$ and K is δ-oriented, if $K = [\sigma, v]$, then σ cannot be an infacet of K. Indeed, the δ-orientation condition gives $v - y/\|v - y\| \in C(f(y), \delta)$ for some $y \in \sigma$ and $f(y) \in C(g, \omega)$ which by Lemma 3.3 implies that $v - y/\|v - y\| \in C(g, \lambda)$ which contradicts $(v - y) \cdot g < 0$ since σ lies in the positive halfspace from H_g. Thus, K must have an infacet containing v, since otherwise K would contain an equilibrium point, but we have assumed that the vector field has no equilibria in $|\mathcal{K}|$.

Now let σ be an infacet containing v. By assumption, the projection $\pi_g \sigma$ is contained in the disk $D := B(v, \rho_g(K)) \cap H_g$. Consider the point $p = v - (\rho_g(K)\lambda/\sqrt{1 - \lambda^2})g$. The cone $C(p, g, \lambda)$ intersects H_g in precisely the disk D, which contains the projection $\pi_g \sigma$. Thus $C(p, g, \lambda)$ contains σ. Since K is δ-oriented and σ is an infacet, the same argument used in the proof of Lemma 3.5 implies that $K \subset C(p, g, \lambda)$. ∎

3.2 Approximating Chain Recurrence

In the previous subsection, we proved local results that apply on regions of the phase space where the direction of the vector field is contained in a cone. These results assert that forward images of the multivalued map on an oriented complex are also contained in a cone with a bound on the location of the apex. Now we would like a global theorem for arbitrary vector fields which asserts that map recurrent sets from oriented complexes are contained in the ϵ-chain recurrent set. Our strategy will be to cover trajectories by cones. The first lemma allows us to find sufficient conditions on the simplices of the complex to build a covering of a trajectory by cones for which the apex of each cone lies in the previous cone, cf. Figure 2.

Lemma 3.9 *Let $0 < \lambda < 1$ and $h \in S^{n-1}$. Suppose K is δ-oriented with barycenter b and $g = f(b)$. Let $r = \rho_g(K)\lambda/\sqrt{1-\lambda^2}$. Suppose further that $K \subset C(x, h, \lambda)$ and the distance from x to any point in K is larger than some $\beta > 0$. Then the apex p such that $K \subset C(p, g, \lambda)$, guaranteed by Lemma 3.8, can be chosen such that $p \in C(x, h, \Lambda)$ where*

$$\Lambda(\lambda, \beta, r) = \frac{\lambda\sqrt{\beta^2 - r^2}}{\beta} - \frac{r\sqrt{1-\lambda^2}}{\beta} \tag{10}$$

provided $r < \beta$.

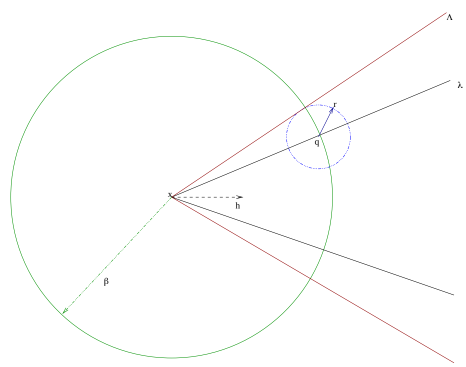

Figure 1: Estimate for $\Lambda(\lambda, \beta, r)$.

Proof: Consider the diagram in Figure 1. We know from Lemma 3.8 that the apex p of the cone $C(p, g, \lambda)$ is within a ball of radius $r = \rho_g(K)\lambda/\sqrt{1-\lambda^2}$ centered at a vertex v of K. We would like to compute the largest value of Λ, i.e. the smallest cone, which contains all balls of radius r centered at a point $q \in C(x, h, \lambda) \cap B(x, \beta)^c$. This occurs when q is on the boundary of the cone $C(x, h, \lambda)$. Let T be a vector through x that is tangent to $B(q, r)$. Then T makes the largest angle, $\cos^{-1}(\Lambda)$, with h when T and h are coplanar with the vector $q - x$. Letting θ represent the angle between the vectors $q - x$ and T, we have that $\cos^{-1}(\Lambda) = \theta + \cos^{-1}(\lambda)$. The result follows from the angle addition formulas and the relations $\sin(\theta) = r/\beta$, $\cos(\theta) = \sqrt{\beta^2 - r^2}/\beta$, and $\sin(\cos^{-1}(\lambda)) = \sqrt{1-\lambda^2}$. ∎

To use the previous lemma, we will need a lower bound on Λ in terms of λ so that $\Lambda \to 1$ as $\lambda \to 1$. In essence this requires a bound on $\rho_g(K)$ for each simplex in terms λ. Choose any function $S(\lambda)$ such that $0 < S(\lambda) < 1/2$ and $S(\lambda) \to 0$ as $\lambda \to 1$ and require that $r < \beta S(\lambda)$, which implies $r < \beta/2$. Then from (10) we can estimate

$$\lambda - S(\lambda)(\lambda + \sqrt{1-\lambda^2}) \leq \Lambda \leq \lambda, \tag{11}$$

providing the necessary lower bound.

We will also need the following theorem on Euler approximations for ordinary differential equations, a proof of which can be found in [4].

Theorem 3.10 *Suppose a vector field f has Lipschitz constant L in some region Ω. Let $x(t), \tilde{x}(t)$ be continuous functions mapping into Ω whose continuous derivatives exist except at possibly finitely many points. If x and $\tilde{x}$ satisfy $x' = f(x)$ with errors ϵ_1 and ϵ_2 respectively for $|t| \leq h$, then*

$$\|x(t) - \tilde{x}(t)\| \leq e^{L|t|}\|x(0) - \tilde{x}(0)\| + \frac{\epsilon_1 + \epsilon_2}{L}(e^{L|t|} - 1) \quad \text{for all } t \in [-h, h].$$

We are now ready to prove the main result of this section which uses the δ-orientation condition to relate certain components of the recurrent set of the multivalued map $\mathbf{F}$ to the ϵ-chain recurrent set.

Definition 3.11 A component $\mathbf{Q}$ of the recurrent set of $\mathbf{F}$ is a *strongly recurrent* if $\mathbf{Q}$ contains at least two polygons which do not share a vertex.

Theorem 3.12 *Let $X \subset \mathbb{R}^n$ be a compact polyhedral set, endowed with a Lipschitz unit vector field $f : X \to S^{n-1}$. For every $\epsilon > 0$ there exists $0 < \lambda < 1$, $\delta > 0$, and $\beta > 0$ such that if $\mathcal{K}$ is a δ-oriented complex on X with $\operatorname{diam}(K) < \beta/2$ and $\rho_g(K) < \beta S(\lambda)\sqrt{1-\lambda^2}/\lambda$ for all $K \in \mathcal{K}$ and $\mathbf{Q}$ is a strongly recurrrent component of the recurrent set of $\mathbf{F} = \mathbf{F}_{\epsilon,1}$ then $|\mathbf{Q}| \subset \mathcal{R}_\epsilon$.*

Proof: Starting from an arbitrary point $x_0 \in |\mathbf{Q}|$, we first construct an Euler path with initial point $p_0 \in B(x_0, \epsilon/2)$ and terminal point $x_1 \in B(\varphi(x_0, \tau), \epsilon) \cap |\mathbf{Q}|$. Iterating this construction yields an ϵ-chain in $|\mathbf{Q}|$ starting at x_0. Then we show that, since $\mathbf{Q}$ is strongly recurrent, this process can be terminated at x_0, so that x_0 is ϵ-chain recurrent.

Step 1: Choices of parameters and scale.

For any choice of scale function $S(\lambda)$ we can bound the right hand side of (9) by

$$\epsilon(\lambda) := \sqrt{2}(1 - \lambda^2 + \lambda^2 S(\lambda) + \sqrt{1-\lambda^2})^{1/2}$$

and $\epsilon(\lambda) \to 0$ as $\lambda \to 1$. Thus choose λ such that $\epsilon(\lambda) < L\epsilon/2(e^{6L} - 1)$, where $L > 0$ is a Lipschitz constant for the vector field.

Consider the function $g(x, y) = f(x) \cdot f(y)$ on $X \times X$. For any $\omega > \lambda$, let $G = g^{-1}(\omega, 1])$. Then G is open in $X \times X$ which implies that for every point $x \in X$ there is a neighborhood U_x of x such that $U_x \times U_x \in G$. Let $\Omega(\omega)$ be the Lebesgue number of the cover $\mathcal{U} = \{U_x : x \in X\}$, and if $\Omega(\omega) > 1$ then set $\Omega(\omega) = 1$. On the neighborhoods in $\mathcal{U}$ the vector field is contained in the cone $C(f(x), \omega)$.

Choose $\beta < \min\{\Omega/6, \epsilon/2e^{6L}\}$ and $\delta \geq \Delta(\omega, \lambda)$.

Step 2: Construction of the Euler approximation.

On the level of simplices, the assumption that $\mathbf{Q}$ is strongly recurrent can be restated as follows. Let $x_0 \in K_0 \subset P$ be an n-simplex and a polygon containing x_0 respectively. Then

P is recurrent which implies that there exist simplices $\{K_i\}_{i=0}^m$ in $\mathcal{K}^{(n)}$ such that $K_0 = K_m$ (assume m is minimal) and $K_{i-1} \cap K_i = \sigma_i$ is an infacet of K_i as in the proof of Corollary 3.6. We will extend this sequence of simplices periodically to get $J = \{K_i\}_{i \geq 0}$. Moreover, since $\mathbf{Q}$ is strongly recurrent we can assume without loss of generality, that $\{K_i\}_{i=0}^m$ are not all part of the same polygon and that $\{K_i\}_{i=0}^m$ do not all share a common vertex.

For any $0 \leq i_0 \leq n-1$ let b be the barycenter of $K = K_{i_0}$, $g = f(b)$, and p_0 the apex of a cone $C(p_0, g, \lambda)$ containing K which is guaranteed by Lemma 3.8. Then by Lemma 3.5, $C(p_0, g, \lambda)$ also contains $|J| \cap B(p_0, \Omega)$. Moreover, by Lemma 3.7 we have that J cannot be recurrent in $C(p_0, g, \lambda)$.

Let r be defined as in Lemma 3.9. Note that $S(\lambda)$ was chosen so that $r < \beta/2$. Since $r + \operatorname{diam}(K_i) < \beta$ for all i and J cannot be recurrent in $B(p_0, \Omega)$, let K_{i_1} be the first simplex after K_{i_0} such that $\operatorname{dist}(K_{i_1}, p_0) \geq 2\beta$, which implies that $\operatorname{dist}(K_{i_1}, p_0) < 3\beta$. As before let b_1 be the barycenter of K_{i_1} and $g_1 = f(b_1)$. Let p_1 be the apex of a cone $C(p_1, g_1, \lambda)$ containing K_{i_1} which is guaranteed by Lemma 3.8. Then $\beta \leq \|p_1 - p_0\| \leq 5\beta$ since $r < \beta$ and $\operatorname{diam}(K_{i_1}) < \beta$.

Iterating this construction yields a sequence of points p_k such that $p_{k+1} \in C(p_k, g_k, \Lambda)$. Let $t_0 = 0$, $t_{k+1} = t_k + \|p_{k+1} - p_k\|$ and $\nu_k = (p_{k+1} - p_k)/\|p_{k+1} - p_k\|$. This defines a piecewise linear approximate trajectory

$$P(t) = p_k + (t - t_k)\nu_k \quad \text{for } t \in [t_k, t_{k+1}]. \tag{12}$$

By Lemma 3.4, we have $\|P'(t) - f(P(t))\| < \epsilon(\lambda)$ for all but finitely many times t for which $P'(t)$ is undefined.

Since $\beta < t_{k+1} - t_k < 5\beta$, iterate the above construction n times where n is the smallest integer such that $n\beta \geq 1$. Then the total time $\tau = \sum_{i=1}^n t_i \in [1, 5 + 5\beta]$. Finally we choose p_{n+1} to be the barycenter of K_{i_n} and the distance $\|p_{n+1} - p_n\|$ is at most $r + \operatorname{diam}(K_{i_n}) < \beta$. Since $p_{n+1} \in C(p_n, g_n, \lambda)$, this final step in the path is still an $\epsilon(\lambda)$-approximate trajectory with total time $1 \leq \tau < 5 + 6\beta < 5 + 6\Omega(\omega)/6 < 6$. This part of the proof is illustrated in Figure 2.

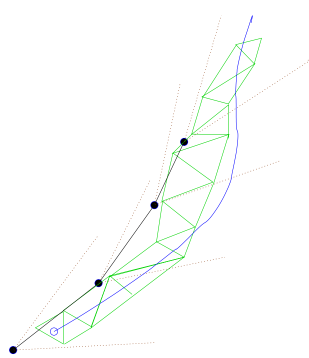

Figure 2: Construction of an Euler path using succesive cone apices.

Step 3: Construction of an ϵ-chain.

Choose any $x_0 \in K_0$. Let $H(t) := \|P(t) - \varphi(x_0, t)\|$. From Theorem 3.10,

$$H(t) \leq e^{Lt} H(0) + \frac{\epsilon(\lambda)(e^{Lt} - 1)}{L}. \tag{13}$$

Now $H(0) = \|x_0 - p_0\| \leq r + \operatorname{diam}(K) < \beta$ and from the choice of β above we see that the first term in the right hand side of (13) is less than $\epsilon/2$. Our choice of λ and the bound on $\epsilon(\lambda)$ imply that the second term in the right hand side of (13) is less than $\epsilon/2$. It follows that $H(t) < \epsilon$ for $0 \leq t < \tau$. Therefore, $x_1 = p_{n+1}$ is in $|J|$ and is at most distance ϵ from x_0. Iterating this construction yields an ϵ-chain $\{x_i\}_{i=0}^m$.

Finally we show that this process can be terminated at $x_m = x_0$ so that x_0 is ϵ-chain recurrent, and we will have shown that $|J| \subset \mathcal{R}_\epsilon$. Proceed in the above construction until the ϵ-chain $x_0, \ldots, x_m$ satisfies the following properties: $m > 2$ and in the construction of the cone apices p_k starting from $x_m \in K_{i_m}$ it happens that $K_0 \subset |J| \cap B(p_k, \Omega) \subset C(p_k, g_k, \lambda)$ but $k\beta < 1$. Let $p_{k+1} = x_0$. Then the path $P(t)$ defined in (12) satisfies the above estimates except that τ_m may be less than 1. However, $\tau_m \leq 6$. So $\{(x_0, \tau_0), \ldots, (x_{m-1}, \tau_{m-1}), (x_m, \tau_m)\}$ is almost an $(\epsilon, 1)$-chain which makes x_0 be ϵ-chain recurrent except it may happen that $\tau_m < 1$.

If so, we define $\tau^*_{m-1} = \tau_{m-1} + \tau_m$ and show $\{(x_0, \tau_0), \ldots, (x_{m-1}, \tau^*_{m-1})\}$ is an $(\epsilon, 1)$-chain with $x_0 \in B(\varphi(x_{m-1}, \tau^*_{m-1}), \epsilon)$ so that x_0 is really ϵ-chain recurrent. We know $\|\varphi(x_{m-1}, \tau_{m-1}) - x_m\| < \epsilon$, and hence Gronwall's inequality implies $\|\varphi(x_{m-1}, \tau^*_{m-1}) - \varphi(x_m, \tau_m)\| < \epsilon e^{6L}$. Since $\|\varphi(x_m, \tau_m) - x_0\| < \epsilon$, we have $\|\varphi(x_{m-1}, \tau^*_{m-1}) - x_0\| < \epsilon(1 + e^{6L})$. Here the original definitions of λ, δ, and β may need to be revised to accommodate $\epsilon(1 + e^{6L})$ in place of ϵ. ∎

A few comments about this theorem are in order. The hypothesis that the component must be strongly recurrent is necessary in dimensions higher than two to avoid situations in which the vector field circulates weakly around a central axis but is not ϵ-chain recurrent. Such components are easily identified in practice since they share a common vertex, and this problem does not arise in two-dimensional flows. Also, the theorem applies only to unit vector fields. However, in practice, one can compute isolated equilibria and consider regions away from equilibria where the vector field can be normalized by a rescaling of time.

4 Examples

In this paper we have established a framework which, given a triangulated region X and a vector field on X, yields an index filtration for a Morse decomposition of the flow on X. We have also given a local criterion on the individual simplices of such a triangulation which guarantees that the resulting index filtration approximates the flow arbitrarily closely. However, the existence of such triangulations is an open problem.

From a computational point of view, one does not necessarily need triangulations which approximate the flow arbitrarily closely. The Conley index can be computed from any isolating block. In this section, we provide some examples of triangulations for various two and three-dimensional flows and the index information obtained from them. The specific details of the numerical algorithm used to compute these triangulations is beyond the scope of this paper and will be addressed in future work. Here our goal is to convince the reader that, even though there is still much work to be done before a general, practical implementation is available for higher-dimensional flows, it is not an unreasonable goal. Moreover, using interval arithmetic techniques, one can rigorously check the transversality conditions for an isolating block and provide computer-assisted proofs.

Example 4.1 (Reverse Van der Pol)

$$\begin{aligned}\dot{x} &= -y\\ \dot{y} &= (x^2-1)y+x.\end{aligned}$$

Figure 3 shows the minimal flow transverse decomposition of a triangulation containing $20,000$ vertices, $39,9940$ triangles, and $27,552$ polygons. From this triangulation, we can extract a component of the recurrent set of the multivalued map which contains the periodic orbit, see Figure 4. This set contains $8,093$ triangles in $5,812$ polygons. This triangulation can be refined to approximate the periodic orbit more closely, as in Figure 5 which shows the recurrent set for the multivalued map containing $20,119$ triangles.

Example 4.2

$$\begin{aligned}\dot{x} &= -x(x+1)-z\\ \dot{y} &= y(2+6x-y)+3\left(x+\frac{z}{3}\right)\\ \dot{z} &= z(2-x+5y)\end{aligned}$$

This 3-dimensional system, related to the ground state problem for a system of coupled semilinear Poisson equations with critical exponents, has a connecting orbit between the equilibria $(0,0,0)$ and $(-1,-1,0)$ as the intersection of a 2-dimensional unstable manifold and a 2-dimensional stable manifold. The connecting orbit is also proven in [3] to be a parabola over the line $x=y, z=0$. An isolating neighborhood $(12,326$ simplices) of this connecting orbit is shown in Figure 6.

References

[1] D. S. Broomhead and A. Iserles, editors. *The dynamics of numerics and the numerics of dynamics*, volume 34 of *The Institute of Mathematics and its Applications Conference Series. New Series*, New York, 1992. The Clarendon Press Oxford University Press.

[2] C. Conley. *Isolated Invariant Sets and the Morse Index*. American Mathematical Society, Providence, RI, 1991.

[3] J. Hulsof and R.C.A.M VanderVorst. Asymptotic behavior of ground states. *Proc. Amer. Math. Soc.*, 124:2423–2431, 1996.

[4] W. Hurewicz. *Lectures on Ordinary Differential Equations*. MIT Press, Cambridge, MA, 1958.

[5] Arieh Iserles. Dynamics of numerics. *Bull. Inst. Math. Appl.*, 30(7-8):106–114, 1994.

[6] W. D. Kalies, K. Mischaikow, and R. C. A. M. VanderVorst. An algorithmic approach to chain recurrence. *Found. Comput. Math.*, 5(4):409–449, 2005.

[7] Konstantin Mischaikow and Marian Mrozek. Conley index. In *Handbook of dynamical systems, Vol. 2*, pages 393–460. North-Holland, Amsterdam, 2002.

[8] J.R. Munkres. *Elements of Algebraic Topology*. Perseus Publishing, Cambridge, MA, 1993.

[9] A. M. Stuart and A. R. Humphries. *Dynamical systems and numerical analysis*, volume 2 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 1996.

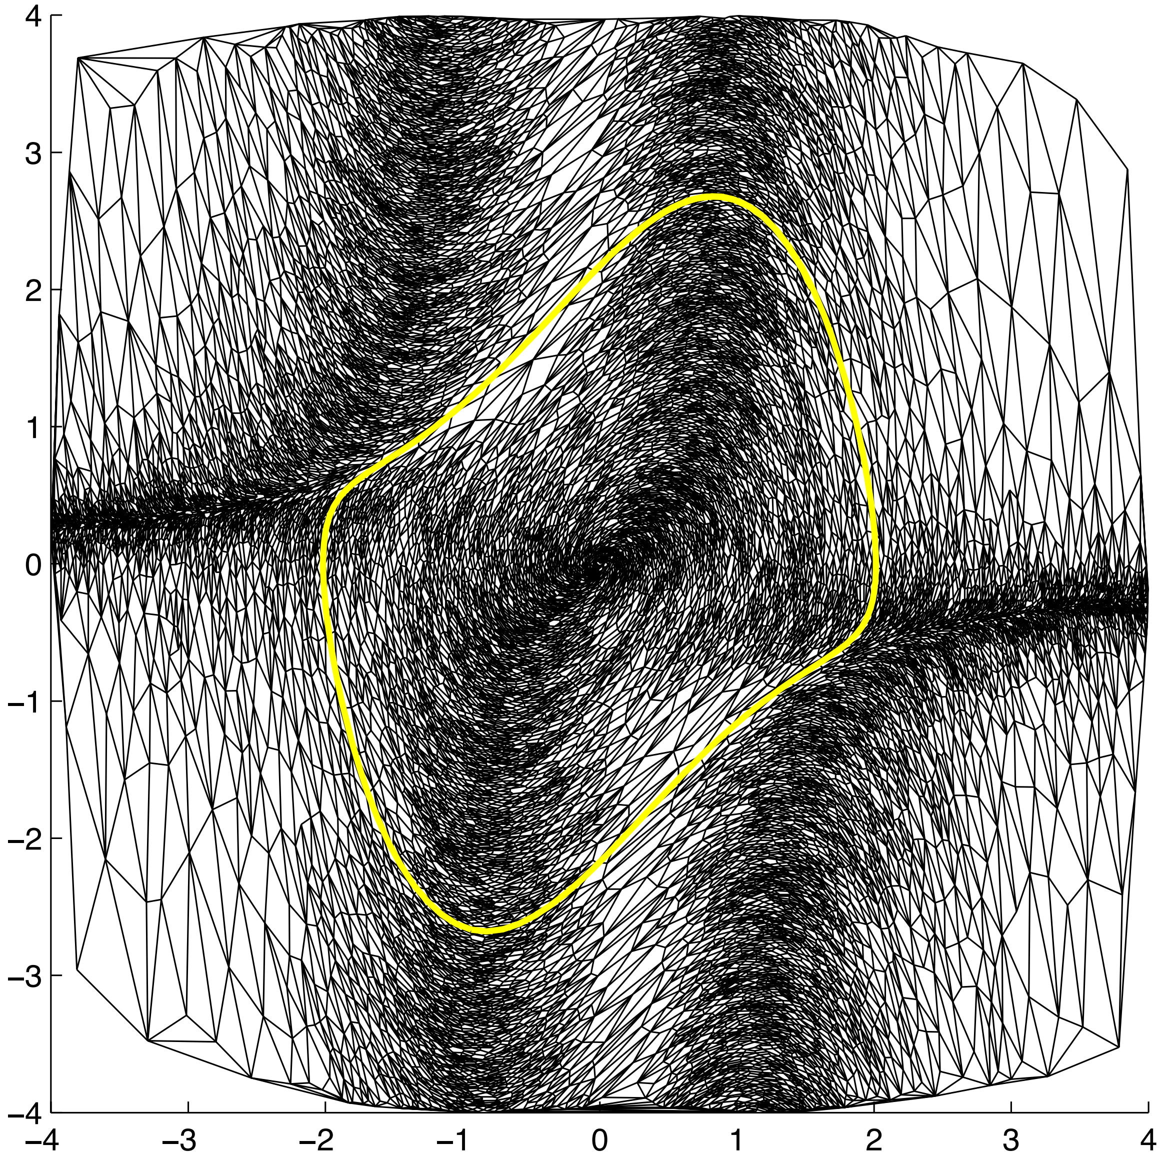

Figure 3: Polygonal decomposition of a triangulation for Reverse Van der Pol containing $20,000$ vertices, $39,9940$ triangles, and $27,552$ polygons. The periodic orbit is shown.

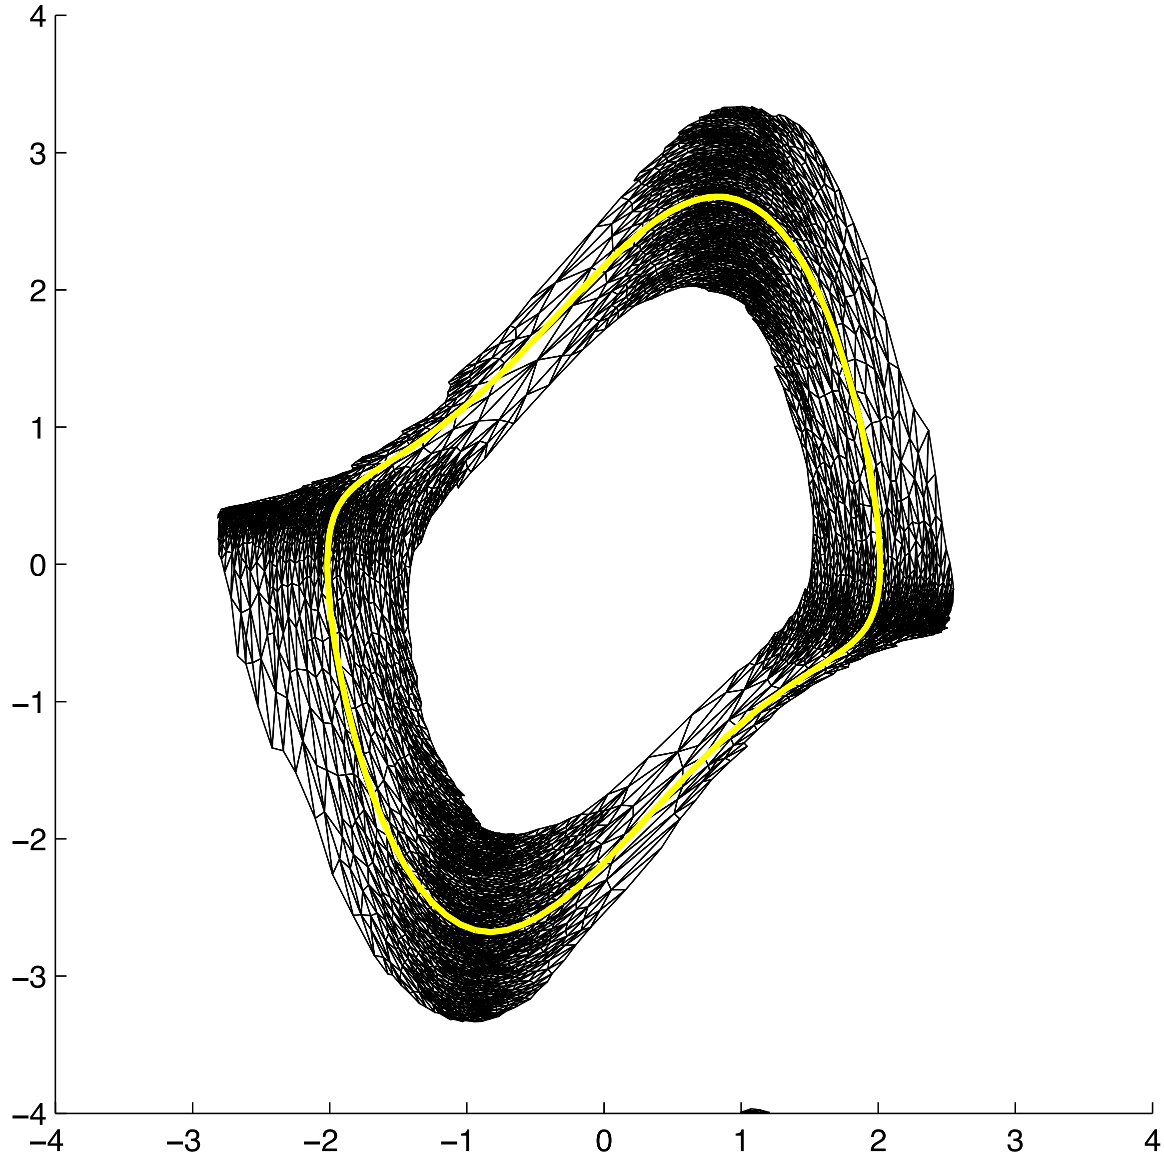

Figure 4: Recurrent component $(8,093$ triangles) of the multivalued map for Reverse Van der Pol.

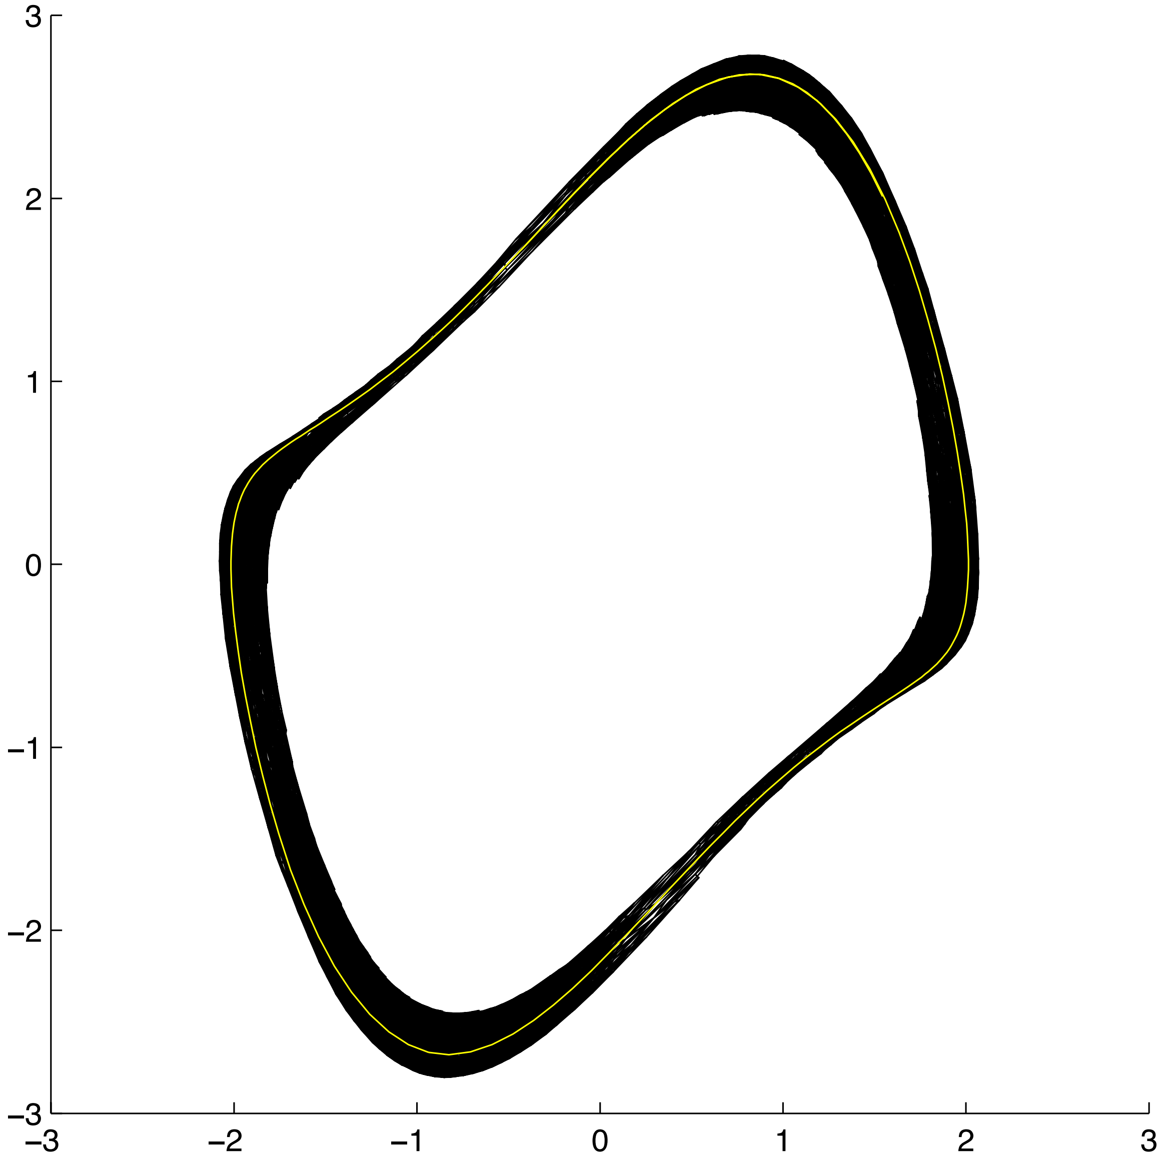

Figure 5: Refinement of the recurrent set $(20, 119$ triangles) of the multivalued map for Reverse Van der Pol.

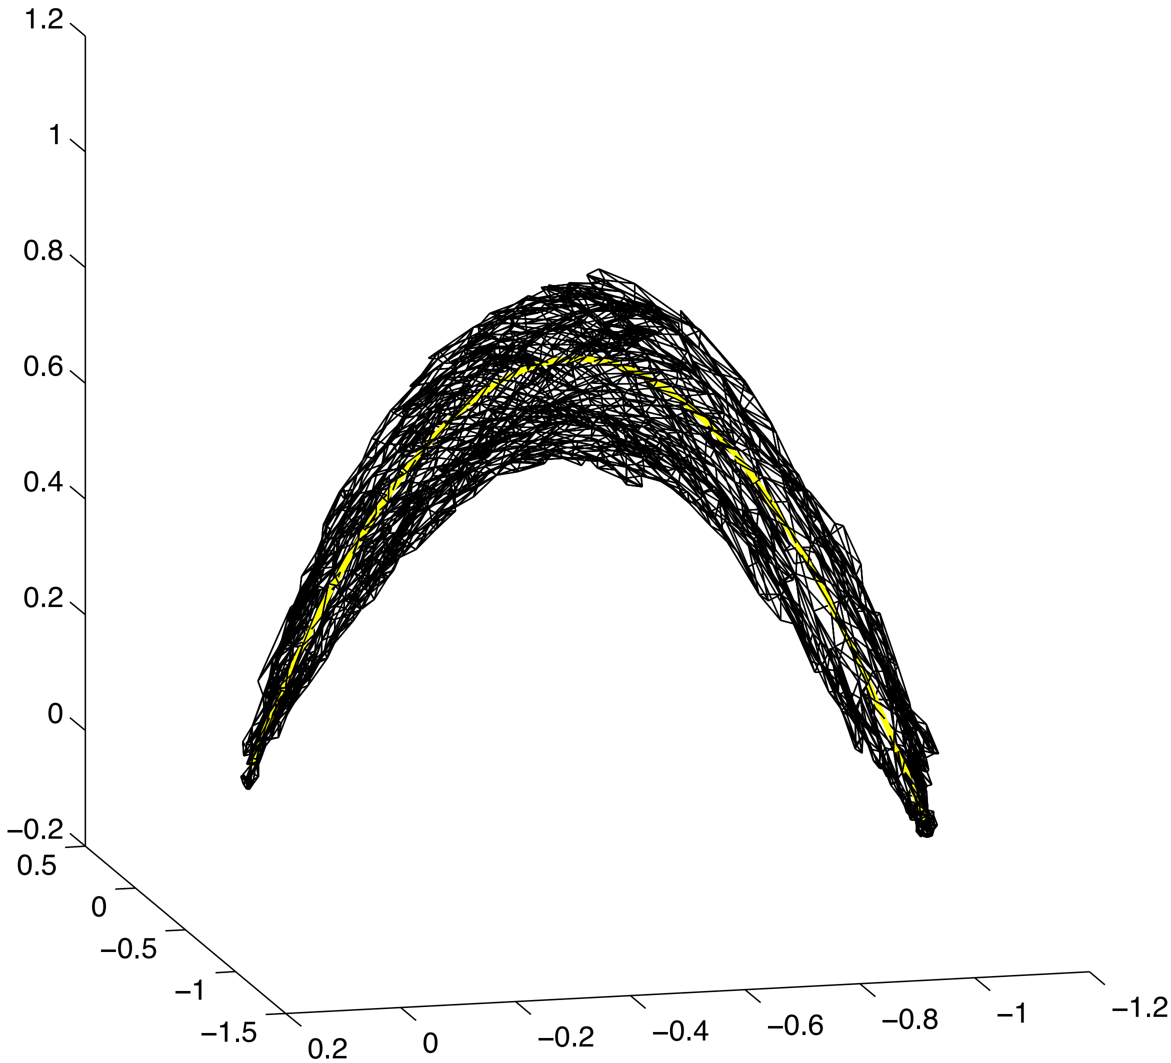

Figure 6: Isolating neighborhood for connecting orbit in Example 4.2 consisting of 12, 326 simplices. Only the boundary edges are displayed.