Reeb Spaces of Piecewise Linear Mappings *

Herbert Edelsbrunner
Dept Computer Science
Duke University, Durham
Raindrop Geomagic, RTP
North Carolina, USA
edels@cs.duke.edu

John Harer
Dept Mathematics
Dept Computer Science
Duke University, Durham
North Carolina, USA
harer@math.duke.edu

Amit K. Patel
Dept Computer Science
Duke University, Durham
North Carolina, USA
akpatel@cs.duke.edu

ABSTRACT

Generalizing the concept of a Reeb graph, the *Reeb space* of a multivariate continuous mapping identifies points of the domain that belong to a common component of the preimage of a point in the range. We study the local and global structure of this space for generic, piecewise linear mappings on a combinatorial manifold.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Computations on discrete structures, Geometrical problems and computations*

General Terms

Algorithms, Theory

Keywords

Smooth and PL topology, combinatorial manifolds, Reeb spaces, cone neighborhoods, triangulations, stratifications, algorithms.

1. INTRODUCTION

This paper advocates Reeb spaces for the structural analysis of continuous, multivariate, scientific data.

Motivation. The current transformation of the physical sciences is driven by the availability of ever larger and more detailed datasets. In many cases, the data samples one or more continuous functions. We model this situation mathematically as a mapping $f : \mathbb{M} \to \mathbb{R}^k$, where $\mathbb{M}$ is the domain and the components of f are the multiple real-valued functions. Commonly asked questions concern the identification of correlated and uncorrelated components and the construction of a small basis that preserves all or most of the information contained in f. The size of the data motivates the extraction of the essential information and their summary. A powerful tool in this context is the *Reeb graph*, which is defined when $k = 1$. It compresses the components of the level sets to points and expresses their relationship by forming a 1-dimensional space. The situation for $k > 1$ is significantly more complicated and the topic of this paper. The central question is how to pack the Reeb graphs of the k components into a single structure. Preferably, this structure reflects the properties of the mapping and is invariant under different choices of basis components. Not surprisingly, the singularities of the mapping play a crucial role, which will become clear when we see how the Jacobi set of the mapping relates to the Reeb space, the proposed single structure.

Beyond summarizing, we are also interested in simplifying the data and its derivatives, in particular the Jacobi set and the Reeb space. While there have been major advances in measuring the size of features [7], the translation of this understanding into effective simplification methods has been slow. Particularly unsuccessful was the attempt to simplify the Jacobi set. Indeed, our original motivation for the study of Reeb spaces was exactly that: to understand the structural constraints that guide the simplification of the Jacobi set. While we feel that this paper is a step in the right direction, we leave the completion of this task to future investigations.

Prior work and results. Because of the importance in visualization, there has been a great deal of work on Reeb graphs of real-valued functions [18]. Motivated by the absence of loops for functions on topologically simple domains, these graphs are sometimes referred to as *contour trees* [1]. The simplification of these trees was discussed in [3]. There is significantly less prior work on the extension of Reeb graphs to Reeb spaces. The existing work is limited to bivariate, generic, smooth mappings.

- Burlet and de Rham study smooth, bivariate mappings on orientable 3-manifolds [2]. Under the assumption that every point of the Jacobi set is definite (appears as a minimum for some linear combination of the two components), they establish relationships between the topology of the 3-manifold and that of the Reeb space. Porto and Furuya extend this work to orientable d-manifolds for $d \geq 3$ [17].
- Motivated by the study of immersions of 3-manifolds in $\mathbb{R}^4$, Levine and coauthors give a complete local classification of points in the Reeb space of bivariate, generic, smooth mappings on orientable as well as non-orientable 3-manifolds [12, 13]. Furuya extends this work to orientable 4-manifolds [9] and Kobayashi and Saeki extend it further to d-manifolds for $d \geq 3$ [11].

In the piecewise linear literature we find only one paper that goes beyond Reeb graphs [6]. Using the Jacobi curves for piecewise linear mappings introduced in [5], it gives a dynamic algorithm for

*The authors are supported in part by DARPA under grants HR0011-05-1-0007 and HR0011-05-0057 and by the NSF under grant DBI-06-06873.

maintaining the Reeb graph in time. The result may be interpreted as sweeping out the Reeb space of a bivariate mapping in which one of the components is time.

In this paper, we consider generic, piecewise linear mappings from a combinatorial d-manifold to $\mathbb{R}^k$. Following the work on generic, smooth mappings, we characterize points of the Reeb space, proving that their neighborhoods are cones over Reeb spaces of one lower dimension. Complementing the local analysis, we show that Reeb spaces have triangulations and coarsest stratifications. Their existence is established constructively. In the case of the triangulation this leads to a polynomial-time algorithm while the construction of the coarsest stratification contains an undecidable subproblem and leads to algorithms only for $k \leq 4$.

Outline. Section 2 provides background from topology. Section 3 introduces Reeb spaces for piecewise linear mappings on combinatorial d-manifolds. Section 4 gives the proof of the local characterization of points of the Reeb space. Section 5 constructs the triangulations and the coarsest stratifications. Section 6 illustrates the results by studying bivariate piecewise linear mappings on orientable 3-manifolds that are generic and simple. Section 7 concludes this paper.

2. BACKGROUND

In this section, we introduce the necessary background on simplicial complexes, piecewise linear functions, and Jacobi sets. For further material on the first two topics we refer the reader to Munkres [16] and, on the last topic, to [5].

Simplicial complexes. An *i-simplex* σ is the convex hull of $i+1$ affinely independent points in some Euclidean space. Letting $u_0, u_1, \ldots, u_i$ be the points, σ is the set of *convex combinations*, that is, points $\sum s_j u_j$ with $\sum s_j = 1$ and $s_j \geq 0$ for all $0 \leq j \leq i$. The *interior* of σ consists of the convex combinations for which all the s_j are strictly positive. The *dimension* of the simplex is $\dim \sigma = i$, which is at most the dimension of the ambient Euclidean space. A *face* of σ is spanned by a non-empty subset of the $i+1$ points. All faces are *proper* except for σ which is an *improper* face of itself. The *boundary* of the simplex, denoted as $\partial\sigma$, consists of all its proper faces. If τ and υ are two disjoint faces of σ with $\dim \tau + \dim \upsilon = \dim \sigma - 1$ then σ is the *join* of the two, $\sigma = \tau * \upsilon$, meaning it is the union of line segments connecting points of τ with points of υ. Any two of these line segments are either equal, disjoint or meet at a common endpoint.

A *simplicial complex* is a finite set of simplices K such that every face of a simplex in K belongs to K and the intersection of any two simplices in K is either empty or a face of both. We call K a *d-complex* and d the *dimension* of K if the largest dimension of any of its simplices is d. The *underlying space* of K is the union of the simplices, $|K| = \bigcup_{\sigma \in K} \sigma$, together with the subspace topology inherited from the ambient space. Avoiding any possible confusion we will sometimes blur the distinction between a complex and its underlying space. A *subcomplex* is a simplicial complex $L \subseteq K$. L is called *full* if it contains every simplex of K whose vertices lie in L. For every non-negative integer $i \leq d$, the *i-skeleton*, denoted as $K^{(i)}$, is the largest subcomplex of dimension i; it consists of all simplices of dimension i or less in K. The 0-skeleton is often referred to as the *vertex set*, $\mathrm{Vert}\, K = K^{(0)}$. The *star* of a simplex σ, denoted as $\mathrm{St}\, \sigma$, is the set of simplices in K that have σ as a face. We get the *closed star* if we add all faces of simplices in the star. The *link* of σ, denoted as $\mathrm{Lk}\, \sigma$, consists of all simplices in the closed star that have an empty intersection with σ. Note that the closed star and the link are complexes while the star is generally not a complex. Extending the notion of underlying space to subsets $L \subseteq K$ we write $|L|$ for the union of interiors of the simplices in the subset. A *subdivision* of K is a simplicial complex with the same underlying space for which every simplex is contained in a simplex in K. Particularly useful is the *barycentric subdivision* which we denote as $\mathrm{Sd}\, K$. To describe it we recall that the *barycenter* of an i-simplex is the average of its $i+1$ vertices. The barycenters of the simplices in K form the vertex set of $\mathrm{Sd}\, K$ and a subset of the barycenters spans a simplex iff the corresponding simplices in K form a chain in which every simplex is a face of the next in the sequence.

We say K *triangulates* a topological space homeomorphic to its underlying space. If K triangulates a d-manifold then every point of $|K|$ has a neighborhood homeomorphic to $\mathbb{R}^d$. However, this does not imply that the link of every i-simplex triangulates a sphere of dimension $d-i-1$. A counterexample to this seemingly plausible property can be found in Edwards [8], see also [19]. We call K a *combinatorial d-manifold* if it satisfies this stronger property, that is, the link of every vertex triangulates the $(d-1)$-sphere and is itself a combinatorial $(d-1)$-manifold. Equivalently, the star of every vertex is isomorphic to the star of a vertex in a subdivision of the d-simplex.

PL mappings and height functions. Let $u_1, u_2, \ldots, u_n$ be the vertices of a simplicial complex K, σ a simplex in K, and a a point of σ. We recall that a is a unique convex combination of the vertices of σ so we can write $a = \sum_{j=1}^n s_j u_j$ with $\sum_{j=1}^n s_j = 1$, $s_j \geq 0$ for all j, and $s_j = 0$ unless u_j is a vertex of σ. The s_j are unique and are called the *barycentric coordinates* of a. We use them to extend a vertex map $\hat{f} : \mathrm{Vert}\, K \to \mathbb{R}^k$ by piecewise linear interpolation to a *piecewise linear* or *PL mapping* $f : |K| \to \mathbb{R}^k$ defined by

$$f(a) \quad = \quad \sum_{j=1}^{n} s_j \hat{f}(u_j).$$

By construction, the restriction of f to a simplex of K is linear. We may think of f as a way to draw K in $\mathbb{R}^k$. Clearly, f is continuous but it is generally not injective. In particular, f restricted to a simplex of dimension beyond k cannot be injective. We call f a *generic PL mapping* if the images of the vertices have no structural properties that can be removed by arbitrarily small perturbations of the vertex map. In particular, we will make use of the following consequence of this general assumption.

I. The restrictions of f to simplices of dimension k or less are injective, that is, the image of every simplex of dimension $i \leq k$ is an i-simplex.

Suppose $h : |K| \to \mathbb{R}$ is a generic PL function on K. By Property I we have $h(u_i) \neq h(u_j)$ whenever u_i and u_j are the two endpoints of an edge in K. We define the *lower link* of a vertex u_j as the collection of simplices in the link whose vertices all have smaller function value than u_j. Symmetrically, the *upper link* is the collection of simplices in the link whose vertices have larger function value:

$$\begin{aligned} \mathrm{Lk}_{-} u_j &= \{\sigma \in \mathrm{Lk}\, u_j \mid a \in \sigma \Rightarrow h(a) < h(u_j)\}; \\ \mathrm{Lk}^{+} u_j &= \{\tau \in \mathrm{Lk}\, u_j \mid a \in \tau \Rightarrow h(a) > h(u_j)\}. \end{aligned}$$

Assuming K is a combinatorial d-manifold, the link of u_j is a triangulation of the $(d-1)$-dimensional sphere, $\mathbb{S}^{d-1}$. The lower and upper links are full subcomplexes of this triangulation. Note that their union is not all of the link as there are simplices that have some of their vertices with higher value, and some with lower value.

We measure the way the lower link is connected using reduced homology with $\mathbb{Z}/2\mathbb{Z}$ coefficients. Following the usual convention, we write $\tilde{\beta}_i$ for the rank of the dimension i reduced homology group. Denoting the ranks of the non-reduced homology groups by β_i we have $\tilde{\beta}_i = \beta_i$ unless $i < 1$. Furthermore $\tilde{\beta}_0 = \beta_0 - 1$ and $\tilde{\beta}_{-1} = 0$ unless the lower link is empty in which case we have $\tilde{\beta}_0 = \beta_0 = 0$ and $\tilde{\beta}_{-1} = 1$. All $\tilde{\beta}_i$ are non-negative integers. We call u_j a *regular vertex* of h if all reduced Betti numbers of its lower link vanish and a *critical vertex*, otherwise. It is a *simple* critical vertex if $\sum \tilde{\beta}_i = 1$. Simple critical points are conveniently classified by the *index* that exceeds the dimension of the non-zero reduced homology group by one; see Table 1. For $d = 3$ it is common to refer to simple critical points of index 0, 1, 2, 3 as *minima*, 1-*saddles*, 2-*saddles*, *maxima*.

type	index	$\tilde{\beta}_{-1}$	$\tilde{\beta}_0$	$\tilde{\beta}_1$	$\tilde{\beta}_2$
regular		0	0	0	0
minimum	0	1	0	0	0
1-saddle	1	0	1	0	0
2-saddle	2	0	0	1	0
maximum	3	0	0	0	1

Table 1: A simple critical point of index i is characterized by $\tilde{\beta}_{i-1} = 1$ and $\tilde{\beta}_j = 0$ for all $j \neq i-1$.

Jacobi sets. We now return to a multivariate generic PL mapping $f : |K| \to \mathbb{R}^k$. Following [5], we consider all linear combinations of the components of f. Equivalently, we let $\vec{u}$ be a unit vector in $\mathbb{S}^{k-1}$ and consider the PL function $h_{\vec{u}} : |K| \to \mathbb{R}$ defined by

$$h_{\vec{u}}(a) = \langle f(a), \vec{u} \rangle,$$

the height of the image of the point a in the direction $\vec{u}$. Assuming $h_{\vec{u}}$ is constant on a simplex τ in K we can define its *lower link* the same way as for a vertex, namely as the collection of simplices in the link whose vertices have function value less than the points of τ. The *upper link* of τ is similarly defined. Assuming the upper and lower links exhaust all vertices of $\mathrm{Lk}\,\tau$, we use the reduced homology of the lower link to decide whether τ is *regular* or *critical* for $h_{\vec{u}}$, and if it is critical whether or not it is *simple*. If τ is a $(k-1)$-simplex then there are exactly two unit vectors for which the height functions they define are constant on τ, namely the unit normals $\vec{u}$ and $-\vec{u}$ of the image of τ in $\mathbb{R}^k$. The lower link of τ under one height function is its upper link under the other, which implies that τ is critical for $h_{\vec{u}}$ iff it is critical for $h_{-\vec{u}}$. In other words, τ has only one chance to be critical. Finally, we define the *Jacobi set* of f as the collection of critical $(k-1)$-simplices together with their faces. These simplices form a subcomplex of K which we denote as J_f.

Property I is needed to unambiguously define the Jacobi set, but it does not imply that this subcomplex has a structure that is as simple as the Jacobi set of a generic smooth mapping. We therefore introduce a second requirement and call $f : |K| \to \mathbb{R}^k$ a *simple, generic PL mapping* if

II. every $(k-1)$-simplex in J_f is a simple critical simplex.

Even property II falls short of implying that the underlying space of J_f is a manifold, but this would be asking too much since even Jacobi sets of generic, smooth mappings are not necessarily manifolds unless k is very small [10].

We will not need the simplicity property until section 6 when we deal with the specific case when K has dimension 3.

3. REEB SPACES

In this section, we introduce the main concept studied in this paper, the Reeb space of a piecewise linear mapping.

Generic preimages. Let K be a combinatorial d-manifold and $f : |K| \to \mathbb{R}^k$ a generic PL mapping. We are interested in the preimage of a point $c \in \mathbb{R}^k$. By Property I, $f^{-1}(c)$ meets every simplex of dimension $i \leq k$ in at most a single point. This implies that it meets every simplex of dimension $k \leq i \leq d$ either in the empty set, a point, or an $(i-k)$-dimensional convex polytope, namely the intersection of the i-simplex with an $(i-k)$-plane contained in the affine hull of the i-simplex. For most points c these polytopes fit together to form a manifold, as we now prove.

GENERIC PREIMAGE LEMMA. Let K be a combinatorial d-manifold, $f : |K| \to \mathbb{R}^k$ a generic PL mapping, and c a point in $\mathbb{R}^k$ not in the image of the $(k-1)$-skeleton. Then $f^{-1}(c)$ is either empty or a manifold of dimension $d-k$.

PROOF. For $d < k$ the preimage of c is empty and for $d = k$ it is either empty or a finite set of points. In both cases there is nothing left to show. We therefore assume $d > k$ for the remainder of this proof.

Let σ be an i-simplex in K. Since K is a combinatorial d-manifold, the link of σ triangulates a sphere of dimension $d-i-1$. Letting u be the barycenter of σ, we construct $B_\sigma = u * |\mathrm{Lk}\,\sigma|$ by drawing a line segment from u to every point in the link. Clearly, B_σ is a PL ball of dimension $d-i$. We further draw a line segment between every point of B_σ and every point of the boundary of σ, as sketched in Figure 1. Any two of these line segments are either

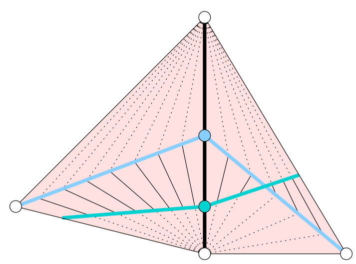

Figure 1: We see a vertical edge and the corresponding 1-ball obtained by connecting its midpoint to the respective third vertices of the two triangles in the star. Connecting every point of the 1-ball to the endpoints of the edge gives a decomposition of the closed star. Pieces of the decomposing line segments form a homeomorphism between the 1-ball and a portion of the preimage $f^{-1}(c)$.

disjoint or meet at a common endpoint, which is either in B_σ or in $\partial\sigma$. Together, the line segments decompose the closed star of σ.

Next we show that for $i = k$ the portion of $f^{-1}(c)$ inside the closed star of σ is homeomorphic to B_σ. Equivalently, every vertex of the preimage has a neighborhood homeomorphic to $\mathbb{R}^{d-k}$. This implies that $f^{-1}(c)$ is indeed a $(d-k)$-manifold. Let σ be a k-simplex in K that contains a point u_σ with $f(u_\sigma) = c$. Because $f^{-1}(c)$ avoids the $(k-1)$-skeleton of K, u_σ belongs to the interior of σ. Let $\tau \in \mathrm{St}\,\sigma$ and let υ be its maximal face disjoint from σ. Hence $\sigma * \upsilon = \tau$ and $u * \upsilon$ is the contribution of τ to B_σ. Letting j be the dimension of τ we have $\dim \upsilon = j - k - 1$ and $\dim(u * \upsilon) = j - k$. Furthermore, $f^{-1}(c)$ intersects τ in a polytope of dimension $j - k$. The line segments in the decomposition of $B_\sigma * \partial\sigma$ define a piecewise linear homeomorphism from $u * \upsilon$

to this polytope; see Figure 1. The collection of such then gives a homeomorphism from B_σ to the intersection of $f^{-1}(c)$ with the closed star of σ. $\square$

Quotient space. Intuitively, the Reeb space of f parameterizes the set of components of preimages of points $c \in \mathbb{R}^k$, we need to describe its topology. By the Generic Preimage Lemma, all but a measure zero subset of these components are manifolds of dimension $d-k$. As we vary c without crossing the image of the $(k-1)$-skeleton these manifolds vary without changing their topological type. Since c has k degrees of freedom this variation has locally the structure of a k-manifold. Only when c belongs to the image of the $(k-1)$-skeleton can we have violations of the manifold property and get shapes that appear as transitions between manifolds of possibly different global connectivity. In summary, we may expect the Reeb space to have the structure of a collection of k-manifolds that are glued to each other in possibly complicated ways. The remainder of this paper show that this is indeed the right intuition. It does this by first formally introducing the Reeb space and then studying its local and global topological properties.

Call two points a and b in $|K|$ *equivalent*, denoted by $a \sim b$, if $f(a) = f(b)$ and a and b belong to the same component of the preimage $f^{-1}(f(a)) = f^{-1}(f(b))$. The *Reeb space* is the quotient space obtained by identifying equivalent points, $\mathbb{W}_f = |K|/\sim$, together with the quotient topology inherited from $|K|$. We already have a map from $|K|$ to $\mathbb{R}^k$, namely f, and one from $|K|$ to $\mathbb{W}_f$, which we call the *quotient map*, q_f. The *Stein factorization* adds another map g from $\mathbb{W}_f$ to $\mathbb{R}^k$ such that the triangle commutes:

$$\begin{array}{ccc} |K| & \xrightarrow{f} & \mathbb{R}^k \\ & {\scriptstyle q_f}\searrow \quad \nearrow{\scriptstyle g} & \\ & \mathbb{W}_f. & \end{array}$$

Since f is generic PL, the dimension of the Reeb space is $d-(d-k) = k$, the same as the image of f. Furthermore, it is not difficult to prove that the Reeb space is Hausdorff, that is, any two different points in $\mathbb{W}_f$ have neighborhoods that are disjoint.

An example. We illustrate the definitions with a mapping from a 3-manifold to the plane. It is convenient to describe the case of a smooth mapping, the extra details that appear in the case of a PL approximation are not difficult. We also note that our example is not compact, so should treated as local since we always assume in the PL case that we have a finite simplicial complex. Consider $f : \mathbb{R}^3 \to \mathbb{R}^2$ defined by its two component functions $f_1(x_1, x_2, x_3) = x_2^3 - x_1 x_2 + x_3^2$ and $f_2(x_1, x_2, x_3) = x_1$. The preimage of a point $c = (s, t)$ is the intersection of two level surfaces, $f_1^{-1}(s) \cap f_2^{-1}(t)$. Setting the two components to s and t we get $x_1 = t$ and $x_3^2 = \gamma(x_2)$ where $\gamma(x_2) = s - x_2^3 + t x_2$. For $t > 0$, γ has a minimum and a maximum, for $t = 0$ it has a single degenerate critical point, and for $t < 0$ it has no critical points. We are interested in the number of roots and in particular the values of x_2 for which $\gamma(x_2) \geq 0$ since only for those do we get a solution to $x_3^2 = \gamma(x_2)$. The odd degree polynomial γ has either 1 or 3 roots except where γ and its derivative have a common zero, which occurs along the fold $27s^2 = 4t^3$. As illustrated in Figure 2, this curve decomposes the (s,t) plane into two regions, γ has three roots above, two roots on, and one root below the fold. Accordingly, $f^{-1}(c)$ has two components above and one component below the fold. It consists of a curve and an isolated point for c on the left branch and of two touching curves for c on the right

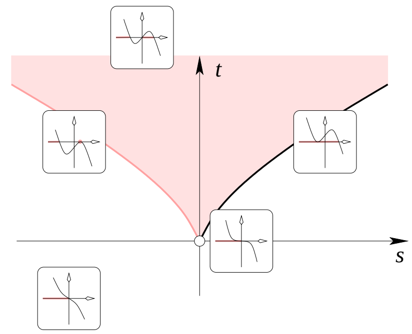

Figure 2: Above the fold the function γ has three roots and thus two intervals in which it is positive. Below the fold γ has one root and only one interval in which it is positive.

branch of the fold. It should be clear how these cases transition between each other as we vary c in the plane. The fold is of course the image of the Jacobi set under the mapping f.

We can re-interpret Figure 2 as a picture of the Reeb space of f. Indeed, it is the image of $\mathbb{W}_f$ under the map g in the Stein factorization. The region above the fold is covered twice and the region below is covered once. Correspondingly, the Reeb space consists of two sheets, one covering the entire plane and the other covering the region above the fold. The latter connects to the former along the right branch of the fold, which is where the two components of $f^{-1}(c)$ come together to merge into a single component. The left branch of the fold is the image of a boundary piece of the second sheet and has nothing to do with the first sheet. In summary, the Reeb space consist of the plane with another two-dimensional sheet attached to it, like a fin sticking out of a fish as in Figure 6, left.

4. LOCAL STRUCTURE

In this section, we prove that every point of the Reeb space has a neighborhood that is a cone over a Reeb space of dimension one less. In stating and proving this result, we follow the work on generic, smooth mappings in [11].

Tubes, cores, and cones. As usual we let K be a combinatorial d-manifold and $f : |K| \to \mathbb{R}^k$ a generic PL mapping. We shed light on the neighborhood structure of a point in the Reeb space by considering its preimage. To be specific, let x be a point of $\mathbb{W}_f$ and $c = g(x)$ its image in $\mathbb{R}^k$. Let B be a closed ball centered at c that is sufficiently small so it intersects the image of a simplex iff this image includes c. Considering the preimages of c and of B, we are interested in the component C of $f^{-1}(c)$ whose image under the quotient map is x and in the component T of $f^{-1}(B)$ that contains C. We call T a *tube* and C its *core*. For example, the core in Figure 5 is a circle and the tube is a solid torus. Define the boundary ∂T of the tube T to be the intersection of T with $f^{-1}(\partial B)$, and let $r = f(a) : \partial T \to \mathbb{R}^k$ be the restriction of f. The corresponding restriction q_r of q_f maps ∂T to $\mathbb{W}_r \subset \mathbb{W}_f$, a subspace whose dimension is $(d-1)-(d-k) = k-1$, one less than the dimension of $\mathbb{W}_f$. Since K is finite, $|K|$ is compact. This implies that ∂T is compact and so is $\mathbb{W}_r$. The *(closed) cone* over $\mathbb{W}_r$ is the space

$$\mathrm{cone}(\mathbb{W}_r) = (\mathbb{W}_r \times [0,1]) / (\mathbb{W}_r \times 1)$$

and its *cone point* is $(\mathbb{W}_r \times 1)/(\mathbb{W}_r \times 1)$. We are now ready to state the first structural result of this paper.

CONE NEIGHBORHOOD THEOREM. Let K be a combinatorial d-manifold, $f : |K| \to \mathbb{R}^k$ a generic PL mapping, and $\mathbb{W}_f$ the Reeb space of f. Then each point $x \in \mathbb{W}_f$ has a homeomorphism from $\mathrm{cone}(\mathbb{W}_r)$ to a closed neighborhood that maps the cone point to x.

To prove this theorem we use that the cone over $\mathbb{W}_r$ is compact and that $\mathbb{W}_f$ is Hausdorff. Every continuous injection from a compact to a Hausdorff space is an embedding [15, page 167]. This means that the compact space and its image are homeomorphic. It thus suffices to construct a continuous injection $\eta : \mathrm{cone}(\mathbb{W}_r) \to \mathbb{W}_f$ that maps the cone point to x. The next paragraph does exactly that.

Constructing an embedding. We begin by constructing the barycentric subdivision of K, slightly modified by placing the new vertices not always at the barycenters of the simplices. The connecting simplices are the same as in the standard definition. Specifically, if C intersects the interior of a simplex σ in K then we choose a point $u_\sigma \in C \cap \mathrm{int}\,\sigma$ as the vertex in $\mathrm{Sd}\,K$ that represents σ. If C does not intersect the interior of σ then we choose a point $u_\sigma = \mathrm{int}\,\sigma - T$, which exists because B is sufficiently small. By construction, there is a subcomplex L of $\mathrm{Sd}\,K$ whose underlying space is the core, $|L| = C$; as illustrated in Figure 3. It is not difficult to prove that L is a full subcomplex. Extending the concept of a star we write $\mathrm{St}\,L$ for the set of simplices in $\mathrm{Sd}\,K$ that have a face in L. The tube is covered in its entirety by the interiors of the simplices in $\mathrm{St}\,L$.

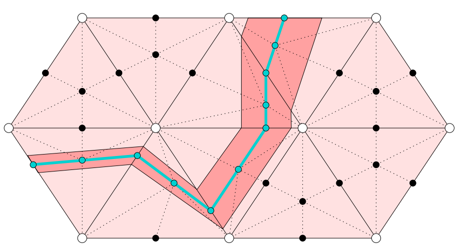

Figure 3: A piece of the barycentric subdivision of K. The white dots mark the vertices of K and the black and shaded dots mark the new vertices of $\mathrm{Sd}\,K$. The shaded path is a piece of the core which is subdivided by the subcomplex L of $\mathrm{Sd}\,K$. The corridor along the path is a piece of the tube which is contained in $\mathrm{St}\,L$.

We first use the barycentric subdivision to establish a continuous map from $\partial T \times [0,1]$ to T whose restriction to $\partial T \times [0,1)$ is a homeomorphism onto $T - C$. Let τ be a simplex in $\mathrm{St}\,L$ but not in L and let σ be the maximal face of τ that belongs to L. We observe that σ is unique because L is full. Let υ be the maximal face of τ that is disjoint from σ and note that $\tau = \sigma * \upsilon$, the join of its two faces. Writing all simplices τ as joins we get a decomposition of the closed star of L into line segments. As usual, any two line segments in this decomposition are either disjoint or meet at a common endpoint. Each point of the core is an endpoint of a collection of line segments. In contrast, a point a of the boundary of the tube belongs to exactly one line segment. Letting b be the endpoint of this line segment in the core we let $\lambda_a : [0,1] \to T$ be the straight line mapping $\lambda_a(t) = (1-t)a + tb$. Combining the maps λ_a over all $a \in \partial T$ gives the map $\lambda : \partial T \times [0,1] \to T$. As anticipated, the restriction of λ to $\partial T \times [0,1) \to T - C$ is a homeomorphism and λ itself is continuous. Finally, define $\gamma : \partial T \times [0,1] \to \mathbb{W}_f$, by setting $\gamma = q_f \circ \lambda$. The new map γ takes $\partial T \times 1$ to the point x. The preimage of every other point y in the image of γ is of the form $U \times t$, where U is the preimage of a point in $\mathbb{W}_r$ and t is in $[0,1)$.

Next we map $\partial T \times [0,1]$ to the cone over $\mathbb{W}_r$. Recall that $r : \partial T \to \mathbb{R}^k$ is the restriction of f to the boundary of the tube and $q_r : \partial T \to \mathbb{W}_r$ is the corresponding restriction of q_f. We extend q_r to a map from $\partial T \times [0,1]$ to $\mathbb{W}_r \times [0,1]$ by taking the product with the identity on the unit interval. Composing this product map with the quotient map $\mathbb{W}_r \times [0,1] \to \mathrm{cone}(\mathbb{W}_r)$, we get $q : \partial T \times [0,1] \to \mathrm{cone}(\mathbb{W}_r)$, which maps $\partial T \times 1$ to the cone point. The preimage of every other point y in $\mathrm{cone}(\mathbb{W}_r)$ is of the form $U \times t$, where U is the preimage of a point in $\mathbb{W}_r$ and t is in $[0,1)$, as before. This finally induces a unique map, η, from the cone to the Reeb space that makes the triangle commute:

$$\begin{tikzcd} & \partial T \times [0,1] \arrow[dl, "q"'] \arrow[dr, "\gamma"] & \\ \mathrm{cone}(\mathbb{W}_r) \arrow[rr, "\eta"] & & \mathbb{W}_f. \end{tikzcd}$$

To finish the proof of the Cone Neighborhood Theorem we just need to realize that η has the two required properties. It is continuous because both γ and q are continuous and it is injective because the preimages of points in the images of q and of γ are the same sets $U \times t$.

5. GLOBAL STRUCTURE

In this section, we show that Reeb spaces have canonical stratifications into manifolds. We give a construction in two steps, first triangulating the Reeb space and second by grouping simplices to form the strata.

Refining arrangement. As before we let K be a combinatorial d-manifold and $f : |K| \to \mathbb{R}^k$ a generic PL function. We also assume $k < d$, otherwise K is itself a triangulation of the Reeb space. To prepare the construction of a triangulation, we refine K by decomposing its simplices into prisms aligned with the fibers of f. Specifically, we take the images of the $(k-1)$-simplices of K in $\mathbb{R}^k$, dissect space with their affine hulls, and decompose the simplices using the preimage of the dissection. By assumption of genericity, the image of every $(k-1)$-simplex $\sigma \in K$ is a $(k-1)$-simplex and its affine hull is a $(k-1)$-dimensional plane in $\mathbb{R}^k$. The collection of such planes dissects $\mathbb{R}^k$ into closed chambers, each a convex polyhedron of dimension k. We call this the *arrangement* defined by the planes [4]. To refine K, we take each simplex τ and decompose it into sets of points that map into a common chamber or a common intersection of chambers. For a k-simplex τ these sets are k-dimensional convex polytopes, the same as the chambers. For a $(k+1)$-simplex τ these sets are $(k+1)$-dimensional prisms each uniquely determined by its top and bottom faces of dimension k. It is allowed that the top and bottom faces touch each other along a common face, generating a partially degenerate prism in between. We show that it is not necessary to study decompositions of simplices of dimension beyond $k+1$.

SKELETON LEMMA. The Reeb space of $f : |K| \to \mathbb{R}^k$ is homeomorphic to the Reeb space of the restriction of f to the $(k+1)$-skeleton of K.

PROOF. Let $e : |K^{(k+1)}| \to \mathbb{R}^k$ be the restriction of f to the $(k+1)$-skeleton and recall that points a and b are equivalent if they map to the same image, $e(a) = e(b) = c \in \mathbb{R}^k$, and belong to the same component of the preimage, $e^{-1}(c)$. By assumption of genericity, this preimage is a complex whose maximal elements are edges, each a line intersecting a $(k+1)$-simplex. In contrast, $f^{-1}(c)$ is a complex whose maximal elements are $(d-k)$-dimensional convex polytopes, each the intersection of a $(d-k)$-dimensional plane with a d-simplex. Since $e^{-1}(c)$ is the 1-skeleton of $f^{-1}(c)$, there is a bijection between the components of the two preimages. Hence there is a bijection between $\mathbb{W}_e$ and $\mathbb{W}_f$. Finally, we observe that the quotient topologies are equivalent implying that the bijection is a homeomorphism between the two Reeb spaces. $\square$

Triangulation. We use the decompositions of the skeleta of K to construct a triangulation of the Reeb space. Let Q be the collection of preimages of chambers decomposing the k-skeleton of K and call two of these polytopes *incident* if they share a common $(k-1)$-dimensional face. Let P be the collection of prisms decomposing the $(k+1)$-skeleton and recall that each $\varphi \in P$ has two k-dimensional faces in Q, its *top face* φ_t and its *bottom face* φ_b. The algorithm partitions Q into blocks, starting with the partition into singletons, $\mathcal{Q} = \{\{\psi\} \mid \psi \in Q\}$. We write Q_ψ for the block that contains ψ.

`for` each prism $\varphi \in P$ `do`
  `if` $Q_{\varphi_t} \neq Q_{\varphi_b}$ `then`
    merge the two blocks into one
  `endif`
`endfor`.

When we merge two blocks we remove both from $\mathcal{Q}$ and add their union as a new block to $\mathcal{Q}$. By construction, all polytopes in a block are preimages of the same chamber in the arrangement. We say two blocks Q_ψ and $Q_{\psi'}$ are *incident* if ψ and ψ' are preimages of different but incident chambers in the arrangement and there are at least two preimages, one in each block, that are incident.

A complex representing the Reeb space of f is readily obtained from the partition into blocks. Specifically, for each block Q_ψ in $\mathcal{Q}$ we take a copy of the chamber $f(\psi)$ and we glue these copies along shared $(k-1)$-faces to reflect the incidence relation among the blocks. We further decompose each polytope into simplices and thus finally get a simplicial complex which we denote as W_f. In summary, we have an algorithm that triangulates the Reeb space of a generic PL mapping f from a combinatorial d-manifold to $\mathbb{R}^k$. Assuming d is a constant, the size of the triangulation and the running time of the algorithm are both polynomial in the size of the combinatorial manifold.

Stratification. In general W_f will be significantly finer than necessary to represent the Reeb space. In a first step towards coarsening the representation, we group simplices to form manifolds. The result will be a *stratification* of $\mathbb{W}_f$, that is, a filtration

$$\emptyset = W^{-1} \subseteq W^0 \subseteq \ldots \subseteq W^k = W_f$$

such that each W^j is a subcomplex of W_f and $S^j = W^j - W^{j-1}$ is either empty or a j-manifold. We call S^j the j-*stratum* of the stratification and each of its components a j-*dimensional piece*. In addition to being a j-manifold, we require that all points of a piece are topologically equivalent. By this we mean that any two points x and y of a piece have closed neighborhoods $N(x)$ and $N(y)$ in $|W_f|$ and a homeomorphism from one to the other that maps x to y and whose restriction to the piece is again a homeomorphism. By the Cone Neighborhood Theorem the closed neighborhoods are cones over $(k-1)$-dimensional Reeb spaces. The requirement of topological equivalence can therefore be reformulated in terms of these spaces. Consider the 2-dimensional Reeb space described in the example of section 3. Its 2-stratum consists of two sheets (mapping to the plane and to the fin), its 1-stratum consists of two curves (mapping to the two branches of the fold), and its 0-stratum is one point (mapping to the origin).

We construct the stratification in the order of decreasing dimension. At the top dimension we initialize S^k to the set of k-simplices, each a piece by itself. Then we add simplices of lower dimension effectively merging and enlarging the pieces. For this we use a boolean subroutine DOESBLEND that decides whether or not a simplex fits into a piece or between pieces of the current stratum. We will prove shortly that each iteration starts with a complex W^j of dimension at most j. Following the same pattern as before we can therefore construct the j-stratum of W_f as the top dimension stratum of W^j.

Set $W^k = W_f$;
`for` $j = k$ `downto` 0 `do`
  initialize W^{j-1} to the $(j-1)$-skeleton of W^j
    and S^j to $W^j - W^{j-1}$;
  `for` $i = j-1$ `downto` 0 `do`
    `for` each i-simplex $\zeta \in W^{j-1}$ `do`
      `if` DOESBLEND(ζ, S^j) `then`
        add ζ to S^j and remove it from W^{j-1}
      `endif`
    `endfor`
  `endfor`
`endfor.`

Note that $S^j = W^j - W^{j-1}$ is maintained throughout the algorithm. We still need to establish that the algorithm constructs what we promise but this depends primarily on the boolean subroutine that decides upon which simplices to add to a stratum.

Recognition. According to the definition of a stratification, we need to satisfy two conditions when we add an i-simplex ζ to the current set S^j, the first guaranteeing that we have a j-manifold and the second that points in the same piece have homeomorphic neighborhoods. We formalize both conditions by considering the second barycentric subdivision and comparing links of vertices in this subdivision. Recall that the first barycentric subdivision contains a vertex $\hat{\xi}$ for each simplex $\xi \in W_f$. We refer to it as a *first generation* vertex of $\mathrm{Sd}^2 W_f = \mathrm{Sd}\,\mathrm{Sd}\,W_f$ noting that all its neighbors are second generation vertices in $\mathrm{Sd}^2 W_f$. The link of $\hat{\xi}$ is a model for the boundary of the closed neighborhood of any point in the interior of ξ. Let $\mathrm{Sd}^2 S^j$ be the subset of simplices in $\mathrm{Sd}^2 W_f$ whose interiors are contained in $|S^j|$. We accept ζ as a new simplex in the j-stratum if the following two conditions are satisfied.

1. The link of $\hat{\zeta}$ in $\mathrm{Sd}^2 S^j$ is a $(j-1)$-sphere.
2. There is a homeomorphism that maps the link of $\hat{\zeta}$ to the link of $\hat{\xi}$ in $\mathrm{Sd}\,W_f$, where ξ is already in S^j and belongs to the star of ζ. We also require that the restriction of this homeomorphism to $\mathrm{Sd}^2 S^j$ is a homeomorphism between the two links which by Condition 1 are both $(j-1)$-spheres.

It is clear that this implementation of the boolean subroutine DOESBLEND maintains S^j as a j-manifold. For the top dimension, $j = k$, this implies that whenever S^k contains a simplex then it also contains the simplices in its star. Symmetrically, whenever $W^{k-1} = W^k - S^k$ contains a simplex it also contains its faces. In other words, W^{k-1} is a complex. We can now use induction over the dimension and prove that W^j is a complex for all j. Similarly, whenever S^j contains a simplex then it also contains its star within W^j. Hence if S^j is non-empty then it is a j-manifold for each j. Finally, we notice that the result of the algorithm does not depend on the order in which the simplices are processed. Indeed, the test of the i-simplex ζ does not depend on whether or not any other i-simplices belong to the j-stratum. We thus have a constructive proof of a global property of the Reeb space.

STRATIFICATION THEOREM. Let K be a combinatorial d-manifold, $k \geq 1$, and $f : |K| \to \mathbb{R}^k$ a generic PL mapping. Then the Reeb space $\mathbb{W}_f$ of f is a stratified space and the W^j as constructed by the algorithm form its coarsest stratification.

We note that the constructive proof is really only an algorithm for $k < 5$. Otherwise, the boolean subroutine attempts to recognize when two triangulated spaces of dimension $k - 1 \geq 4$ are homeomorphic, which is undecidable as proved by Markov [14].

6. THE ORIENTABLE 3-MANIFOLD CASE

We take a closer look at the Reeb spaces for PL mappings from an orientable 3-manifold to the plane. In particular, we give a complete case analysis of local cones that arise for simple generic such mappings. To do this we will finally employ the simplicity condition introduced at the end of section 2.

Genericity and simplicity. Let K be a compact combinatorial orientable 3-manifold without boundary and $f : |K| \to \mathbb{R}^2$ a PL mapping. We assume that f is generic and simple. Specifically, we require that

- I'. the intersection of a level set of f with $|K^{(1)}|$ is empty, one point, or two points each in the interior of an edge;
- II'. the Jacobi set of f is a 1-manifold, that is, each edge of J_f is a simple critical edge and each vertex of J_f is endpoint of exactly two edges in J_f.

Recall that to define the lower link of an edge we use the function $h_{\vec{u}} : |K| \to \mathbb{R}$ which maps a point a to the height of $f(a)$ in the direction $\vec{u} \in \mathbb{S}^1$ normal to the edge. A critical edge is simple iff all reduced Betti numbers of this lower link vanish, except for one which is equal to 1. There are three possibilities, namely $\tilde{\beta}_{-1} = 1$ (the cross-section of the simplex is a minimum), $\tilde{\beta}_0 = 1$ (a saddle), and $\tilde{\beta}_1 = 1$ (a maximum); see Figure 4. Condition II' implies that

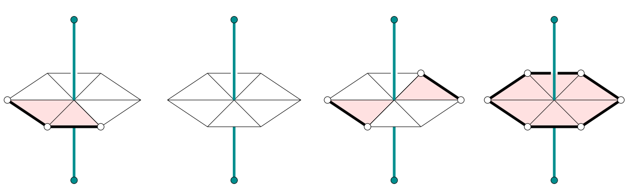

Figure 4: From left to right: a regular edge and three simple critical edges. Each edge is shown with its lower link and a cross-section of its star.

J_f contains no duplicate edges and no duplicate vertices, where by the latter we mean that each endpoint of an edge in J_f belongs to exactly one other edge in J_f; see also [5].

Walks and sheets. To enumerate the cones that may arise, we let $x \in \mathbb{W}_f$ be a point of the Reeb space and B a small closed disk with center $c = g(x)$ in the plane, as in section 4. Furthermore, the core, C, is the component of $f^{-1}(c)$ whose image under q_f is x, and the tube, T, is the component of $f^{-1}(B)$ that contains C. Recall that $r : \partial T \to \mathbb{R}^2$ is the restriction of f that maps the boundary of the tube to the boundary of the disk and that x has a closed neighborhood in $\mathbb{W}_f$ that is homeomorphic to the cone over $\mathbb{W}_r = q_f(\partial T)$. It thus suffices to understand the structure of $\mathbb{W}_r$, which we study by walking the circle in a counter-clockwise order using $\alpha : [0, 1] \to \mathbb{R}^2$ with $\operatorname{im} \alpha = \partial B$. The walk begins and ends at the point $p = \alpha(0) = \alpha(1)$. Letting q be the antipodal point, we also walk the straight diameter using $\beta : [0, 1] \to \mathbb{R}^2$ with $p = \beta(0)$, $c = \beta(\frac{1}{2})$, and $q = \beta(1)$. Note that each point of the two walks is the image of a curve in the tube. The two one-parameter families sweep out the boundary of the tube and another surface we refer to as the *divider*, $D = T \cap f^{-1} \circ \beta[0, 1]$. To describe the two sweeps we define

$$\begin{aligned} \partial T_s &= T \cap f^{-1} \circ \alpha[0, s]; \\ D_s &= T \cap f^{-1} \circ \beta[0, s], \end{aligned}$$

for each $0 \leq s \leq 1$. The simplest of all possible cases is illustrated in Figure 5. There, none of the points in the disk B is critical. Hence, the preimage of every point of $\alpha[0, 1]$ is a closed curve, and the same is true for the preimage of every point of $\beta[0, 1]$. It follows that ∂T_s is an annulus, for every $0 < s < 1$, that closes up to form a torus when s reaches 1. Similarly, D_s is an annulus, for every $0 < s < 1$, and it remains one until the end. The divider, D, is therefore an annulus bounded by two closed curves which it shares with ∂T. Since D also contains the core, C, the only possible configuration is the one depicted in Figure 5. In this particular case,

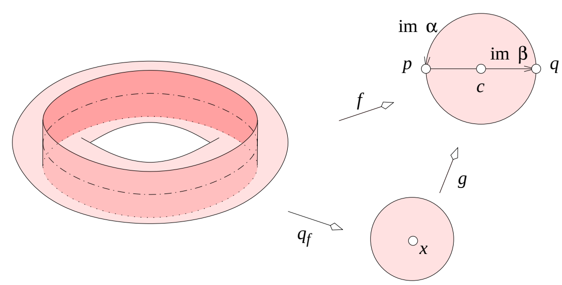

Figure 5: Left: the tube, its core, and the annulus that divides the tube into two. Right: the closed neighborhood of x in the Reeb space and its image in the plane.

the point x belongs to a *sheet*, that is, a piece of the 2-stratum of the Reeb space. The points in the neighborhood correspond to closed curves forming a fibration of the tube.

Choosing the walks. For the more complicated cases it will be convenient to choose the two walks such that p has a connected preimage and both preimages avoid the 1-skeleton of the barycentric subdivision of K. We use the curves sweeping out the divider to prove that such points p and q exist.

ENDPOINT LEMMA. There exist antipodal points p and q of ∂B such that $r^{-1}(p)$ is connected and $r^{-1}(p)$ and $r^{-1}(q)$ both have empty intersection with $(\operatorname{Sd} K)^{(1)}$.

PROOF. Each pair of antipodal points corresponds to a direction $\vec{u} \in \mathbb{S}^1$ such that $\vec{u}$ is a positive multiple of $q - p$. Let $\beta_{\vec{u}} : [0, 1] \to \mathbb{R}^2$ be the corresponding diagonal walk and $D(\vec{u})$ the corresponding divider. Note that the dividers all share the core but are otherwise disjoint.

Fixing a direction $\vec{u}$ and a point a in the core, we consider how the curve $T \cap f^{-1} \circ \beta_{\vec{u}}(s)$, which sweeps out $D(\vec{u})$ as s goes from 0 to 1, intersects a sufficiently small neighborhood $N(a)$ of a in $|K|$. If $a \notin |J_f|$ then the curve looks locally like a line that sweeps over a, passing it at $s = \frac{1}{2}$. Hence $N(a)$ intersects the curve in a connected piece, if at all. If $a \in |J_f|$ then a curve approaches a, pinching off to a single point or recombines leaving a in two different directions. In the former case, we see a closed curve shrinking to a point or the other way round. In the later case, locally we see the usual saddle picture of two pieces that look like the two branches of a hyperbola passing through its pair of asymptotic lines. The two pieces are globally connected along a component of the curve before meeting at a but are not connected after meeting at a, or the other way round. There is an open semi-circle of directions $\vec{u}$ such that $N(a)$ intersects a single component of the curve. This semi-circle is determined by the image of the edge or edges in J_f that contains the point a. By Condition I', there are at most two points in the core that belong to $|J_f|$ and by Condition II' at most two edges in J_f are adjacent to a vertex in the Jacobi set. The corresponding two semi-circles are defined by the images of two different edges in J_f. It follows that the two line segments intersect at c and the corresponding semi-circles intersect in an arc of non-zero length. Picking p on this arc implies that r^{-1} is connected. To satisfy the second requirement of avoiding the 1-skeleton of $\mathrm{Sd}\, K$ we just need to choose p outside a measure zero subset of the arc. $\square$

Arcs. Beyond sheets, the next more complicated case is when the boundary of the tube meets the Jacobi set in two points, a' and a'', in the interior of a single edge or in the interior of two different edges both edges of the same type. The core intersects $|J_f|$ in a single point, a. The points a' and a'' belong to different edges when a is a vertex of J_f and in this case, we assume both edges are definite or both indefinite. This implies two cases and in both the point x belongs to an *arc*, that is, a piece of the 1-stratum. The point a is the sole interior critical point of $\beta^{-1} \circ f : D \to [0, 1]$ and the points a' and a'' the sole interior critical points of $\alpha^{-1} \circ f : \partial T \to [0, 1)$.

CASE A.1. The edge(s) of J_f that contains a' and a'' is(are) definite. The tube is a ball obtained by thickening the point a. The divider, D, depends on the choice of the diameter since the preimage of q may be empty or one curve. Walking along the circle, we start with a single curve that shrinks to the point a' and then reappears from the point a'' sweeping out a sphere. The Reeb space is locally a half-plane, like at a point on the left branch of the fold in Figure 2.

CASE A.2. The edge(s) of J_f that contains a' and a'' is(are) indefinite. The tube is a solid double torus obtained by thickening the figure-8 curve that crosses itself at a. The divider, D, depends on the choice of the diameter since the preimage of q may consist of one or two curves. Walking around the circle, we start with a single curve that splits into two at a' that later merges at a'' to form again a single curve. The Reeb space is locally a book with three pages, like at a point on the right branch of the fold in Figure 2.

Nodes. Next we consider the case in which the core meets J_f at a vertex, a, one of the incident edges of J_f is definite and the other incident edge is indefinite. By Condition II', the boundary of the tube meets the Jacobi set in two points a' and a''. The point a is the sole interior critical point of $\beta^{-1} \circ f : D \to [0, 1]$ and the points a' and a'' are the sole interior critical points of $\alpha^{-1} \circ f : \partial T \to [0, 1)$.

CASE N.1. Assuming a' belongs to the indefinite and a'' to the definite edge, the walk around the circle starts with a single curve that splits into two at a' of which one shrinks to a point at a''. The divider, D, depends on the choice of the diameter. Specifically, the preimage of the endpoint q may consist of one or of two curves. In the former case we have one curve that persists along the entire diameter and in the latter case we start with one curve and get another expanding around a. In either case, the core is a single curve and the Reeb space is locally a disk with a fin sticking out, like at the origin in Figure 2; see also Figure 6, left. The point x is a *node* of the Reeb space, that is, a piece of the 0-stratum.

In the most complicated case the core meets the Jacobi set in two points, a and b. Each of the two points lies in the interior of an indefinite edge, else the core would be disconnected. The boundary of the tube meets the Jacobi set in four points, a', a'', b', and b''. Here, a and b are the sole interior critical points of $\beta^{-1} \circ f : D \to [0, 1]$ and a', a'', b', b'' are the sole interior critical points of $\alpha^{-1} \circ f : \partial T \to [0, 1)$. There are two cases and in both the point x is a node.

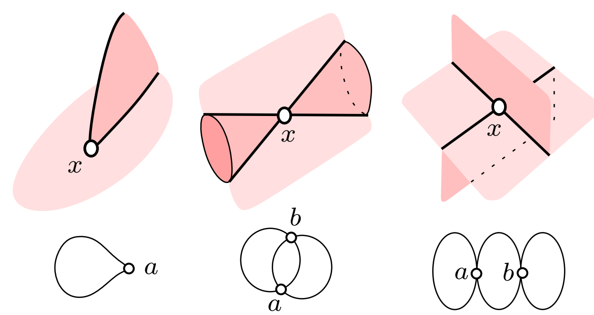

Figure 6: From left to right: the local cone and the core for the Cases N.1, N.2, N.3 at which the point x is a node of the Reeb space.

CASE N.2. Walking along the diameter, we start with a single curve that gets pinched at a and at b with the net effect that it remains a single curve. Knowing that the preimage of q is a single curve determines the boundary of the tube. Using α to walk the circle, we start with a single curve that splits into two curves at a' and then merges into a single curve at b'. As we continue, the single curve splits into two curves at a'' and once again it merges into a single curve at b''. The core consists of two circles that meet at two points, a and b, and the tube is a solid triple torus; see Figure 6, middle.

CASE N.3. Walking along the diameter, we start with a single curve that splits into three curves at a and b. Knowing that the preimage of q consists of three curves again determines the boundary of the tube. As we walk along the circle, the single curve splits into two at a', one of the two splits into two at b', giving a total of three curves. As we continue, two of the three curves merge at a'' and the remaining two merge at b''. The core is a double figure-8 and the tube a solid triple torus; see Figure 6, right.

7. DISCUSSION

The main contribution of this paper is the introduction of Reeb spaces for multivariate, piecewise linear mappings on combinatorial manifolds and the analysis of their local and global structure. There are several open questions that remain.

- The structural assumption that the domain is a combinatorial manifold is critically used in the proof of the Generic Preimage Lemma. Which of our results are not true for general triangulations of manifolds and which do extend? What can be said about Reeb spaces of mappings on simplicial complexes that do not triangulate manifolds?
- How can we effectively simplify the Reeb space of a multivariate mapping? How does this simplification interact with the simplification of the Jacobi set?

Acknowledgments

The third author thanks James Damon for references to work on Reeb spaces of smooth functions. He also thanks Ajith Mascarenhas and Valerio Pascucci for helpful discussions.

8. REFERENCES

[1] M. VAN KREVELD, R. VAN OOSTRUM, C. L. BAJAJ, V. PASCUCCI AND D. R. SCHIKORE. Contour trees and small seed sets for isosurface traversal. *In* "Proc. 13th Ann. Sympos. Comput. Geom., 1997", 212–220.

[2] O. BURLET AND G. DE RHAM. Sur certaines applications génériques d'une variété close à 3 dimensions dans le plan. *L'Enseign. Math.* **20** (1974), 275–292.

[3] H. CARR, J. SNOEYINK AND M. VAN DE PANNE. Simplifying flexible iso-surfaces using local geometric measures. *In* "Proc. IEEE Conf. Visualization, 2004", 497–504.

[4] H. EDELSBRUNNER. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Heidelberg, Germany, 1987.

[5] H. EDELSBRUNNER AND J. HARER. Jacobi sets of multiple Morse functions. In F. Cucker, R. Devore and P. Olver (eds.), *Foundations of Computational Mathematics, Minneapolis 2002*, 35–57, Cambridge University Press, 2004.

[6] H. EDELSBRUNNER, J. HARER, A. MASCARENHAS, V. PASCUCCI AND J. SNOEYINK. Time-varying Reeb graphs for continuous space-time data. *Comput. Geom. Theory Appl.*, to appear.

[7] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.* **28** (2002), 511–533.

[8] R. D. EDWARDS. Approximating certain cell-like maps by homeomorphisms. *Notices Amer. Math. Soc.* **24** (1977), A649.

[9] Y. K. S. FURUYA. Sobre applicações genéricas $\mathbb{M}^4 \rightarrow \mathbb{R}^2$. Tese, Departamento de Matemática, ICMSC-USP, São Carlos, Brazil, 1986.

[10] M. GOLUBITSKY AND V. GUILLEMIN. *Stable Mappings and Their Singularities.* Springer-Verlag, New York, 1973.

[11] M. KOBAYASHI AND O. SAEKI. Simplifying stable mappings into the plane from a global viewpoint. *Trans. Amer. Math. Soc.* **348** (1996), 2607–2636.

[12] L. KUSHNER, H. LEVINE AND J. P. PORTO. Mapping three-manifolds into the plane I. *Bol. Soc. Mat. Mexicana* **29** (1984), 11–33.

[13] H. LEVINE. *Classifying Immersions into $\mathbb{R}^4$ over Stable Maps of 3-Manifolds in $\mathbb{R}^2$*. Springer-Verlag, Berlin, Germany, Lecture Notes in Mathematics **1157**, 1985.

[14] A. A. MARKOV. Insolubility of the problem of homeomorphy. In *Proc. Int. Congr. Math.*, 1958, 14–21.

[15] J. R. MUNKRES. *Topology*. Second edition, Prentice Hall, Englewood Cliffs, New Jersey, 2000.

[16] J. R. MUNKRES. *Elements of Algebraic Topology.* Addison-Wesley, Redwood City, California, 1984.

[17] J. P. PORTO AND Y. K. S. FURUYA. On special generic maps from a closed manifold into the plane. *Topology Appl.* **35** (1990), 41–52.

[18] G. REEB. Sur les points singuliers d'une forme de Pfaff complèment intégrable ou d'une fonction numérique. *Comptes Rendus de L'Académie ses Séances, Paris* **222** (1946), 847–849.

[19] D. P. SULLIVAN. On the Hauptvermutung for manifolds. *Bull. Amer. Math. Soc.* **73** (1967), 598–600.