

Morse Complexes for Piecewise Linear 3-Manifolds *

Herbert Edelsbrunner[†], John Harer[‡], Vijay Natarajan[§] and Valerio Pascucci[¶]

Abstract

We define the Morse complex of a Morse function over a 3-manifold as the overlay of the stable and unstable manifolds of all critical points. In the generic case, its 3-dimensional cells are shaped like crystals and are separated by quadrangular faces. In this paper, we give a combinatorial algorithm for constructing such complexes for piecewise linear data.

Keywords. Computational geometry and topology, densities, iso-surfaces, Morse theory, gradient field.

1 Introduction

Morse functions are used in differential topology as a tool to study the topology of manifolds [11, 12]. We use their results but pursue a different goal, namely that of studying topological features in natural phenomena.

Motivation. A three-dimensional Morse function is a generic smooth map from a 3-manifold to the real numbers. There is an abundance of natural phenomena that can be modeled by such functions. In oceanography, we study the distribution of temperature and other measurements over the Earth's oceans. In medical imaging, we reconstruct the inside of a living body from density distributions measured by MRI and other sensing technology. In x-ray crystallography, we determine the conformations of proteins and other

molecules from electron densities derived from x-ray diffractions. In each case, essential information is obtained from variations of the density over the space. Morse theory offers the basic mathematical language to reason qualitatively and quantitatively about this variation. In oceanography, we might be interested in the temperature extrema and how they change over time. In medical imaging, we use sharp changes in density to segment the data into organs and other functional body parts. In x-ray crystallography, we reconstruct geometric structure by following ridges connecting maxima in electron density.

Prior work on related structures. Three-dimensional densities are commonly visualized by drawing one or several iso-surfaces, which are sets of constant density value. Generically, such a set is a 2-manifold that divides the space into two domains. The 1-parameter family of iso-surfaces sweeps out each cell in the Morse complex in a predictable manner, starting at the minimum and proceeding towards the opposite maximum while crossing the boundary everywhere at a right angle. The most popular method for computing an iso-surface is the marching cubes algorithm, which assumes the density is given by its values at the vertices of a regular grid [10]. Extensions and improvements of this algorithm can be found in [8, 22].

The main performance limitation of the marching cubes algorithm is due to the brute-force traversal of all the cells in the mesh. Faster algorithms avoid traversing empty regions and trace the components of one isosurface by local searches that start from a small set of seed edges [2]. A minimal collection of seeds for all possible iso-surfaces is provided by a minimal covering of the Reeb graph. This graph is a tree, called the contour tree, for simply connected domains [21], but can have loops for more general 3-manifolds [16]. The Reeb graph keeps track of the number of components of the iso-surfaces, but has no geometric information that relates to the gradient flow, as contained in the Morse complex. Extensions and improvements of the original algorithm for computing contour trees can be found in [3, 20, 15].

Another concept related to Morse complexes is the medial axis of a shape in three-dimensional Euclidean space.

*Research by the first and third authors is partially supported by NSF under grants CCR-97-12088, EIA-99-72879, and CCR-00-86013. Research by the second author is partially supported by NSF under grant DMS-01-07621. This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

[†]Department of Computer Science, Duke University, Durham, North Carolina, Raindrop Geomagic, Research Triangle Park, North Carolina, and Lawrence Livermore National Laboratory, Livermore, California.

[‡]Department of Mathematics, Duke University, Durham, North Carolina.

[§]Department of Computer Science, Duke University, Durham, and Lawrence Livermore National Laboratory, Livermore, California.

[¶]Lawrence Livermore National Laboratory, Livermore, California.

The medial axis transform was first introduced in [1] and defined as the set of centers of spheres that touch the boundary of the shape in at least two points but do not cross it. Medial axes are used in a wide variety of applications including shape representation [17, 4], modeling [19], mesh generation [18], image processing [14], computer vision [23], and motion planning [9]. If we define the signed distance from the boundary as the density, we may create the medial axis by stitching together descending 1- and 2-manifolds on one side of the boundary and ascending 1- and 2-manifolds on the other.

Results. A fundamental difficulty in applying Morse theoretic ideas is the lack of smoothness in real data. Most commonly, information is gathered by point probes, and to turn these probes into a generic smooth function is a formidable task. We argue that the construction of such a function is also a questionable step if the goal is to compute and study topological features in the data, mostly because understanding the latter seems necessary to successfully do the former. Instead, we take a combinatorial approach and simulate smoothness to the extent necessary to make things work. The main results of this paper are combinatorial and algorithmic in nature:

- (i) the introduction of quasi Morse complexes as combinatorial analogs of the CW complexes defined by the stable and unstable manifolds of smooth gradient fields;
- (ii) a combinatorial algorithm for constructing a quasi Morse complex with guaranteed structural correctness.

We believe that these results lay the ground-work for a large-scale application of Morse theoretic ideas to data sets in the sciences, engineering and medicine.

Outline. Section 2 presents the background from Morse theory we need in this paper. Section 3 introduces additional background from combinatorial topology and lays out our ideas about mimicking smooth structure for piecewise linear maps. Sections 4, 5, and 6 describe the algorithm for constructing a quasi Morse complex for three-dimensional piecewise linear density data. Section 7 concludes the paper.

2 Smooth 3-Manifolds

In this section, we introduce the Morse theoretic concepts used in this paper. We refer to [11, 12] for further background.

Morse functions. Let M be a smooth compact 3-manifold without boundary. Examples are the 3-sphere, which consists of all points at unit distance from the origin in \mathbb{R}^4 , and the 3-torus, which is obtained by identifying opposite squares of a three-dimensional cube. Let $f : M \rightarrow \mathbb{R}$ be a

smooth map. The differential of f at a point $p \in M$ is a linear map from the tangent space at p to \mathbb{R} , $df_p : TM_p \rightarrow \mathbb{R}$. The point p is *critical* if df_p is the zero map, otherwise it is *regular*. Given a local coordinate system, the *Hessian* at a critical point is the matrix of second order partial derivatives,

$$H(p) = \begin{bmatrix} \frac{\partial^2 f}{\partial x_1^2}(p) & \frac{\partial^2 f}{\partial x_1 \partial x_2}(p) & \frac{\partial^2 f}{\partial x_1 \partial x_3}(p) \\ \frac{\partial^2 f}{\partial x_2 \partial x_1}(p) & \frac{\partial^2 f}{\partial x_2^2}(p) & \frac{\partial^2 f}{\partial x_2 \partial x_3}(p) \\ \frac{\partial^2 f}{\partial x_3 \partial x_1}(p) & \frac{\partial^2 f}{\partial x_3 \partial x_2}(p) & \frac{\partial^2 f}{\partial x_3^2}(p) \end{bmatrix}.$$

The critical point p is *non-degenerate* if the Hessian is non-singular. The function f is called a *Morse function* if all critical points are non-degenerate and $f(p) \neq f(q)$ whenever $p \neq q$ are critical. The Morse Lemma states that if p is non-degenerate we can choose local coordinates such that

$$f(x_1, x_2, x_3) = f(p) \pm x_1^2 \pm x_2^2 \pm x_3^2,$$

for some choice of signs. Note this implies that non-degenerate critical points are isolated. The number of minuses is the *index* of f at the critical point p . It is independent of the coordinate system and equals the number of negative eigenvalues of $H(p)$.

In three dimensions, there are four types of non-degenerate critical points: *maxima* have index 3, *2-saddles* have index 2, *1-saddles* have index 1, and *minima* have index 0. We get intuitive local pictures by drawing a small sphere around a critical point and marking the points with $f(x) < f(p)$ as *ocean*. The ocean on the sphere around a maximum is the entire sphere, and on the sphere around a minimum is the empty set. Figure 1 illustrates the local pictures of the two remaining types of critical points, along with that of a regular point.




Figure 1: The ocean on the sphere around a 2-saddle, a 1-saddle, and a regular point is a belt, the two poles, and a hemi-sphere.

Descending and ascending manifolds. Given a Riemannian metric on M and a local coordinate system with orthonormal tangent vectors $\frac{\partial}{\partial x_i}(p)$, the *gradient* of f at p is

$$\nabla f(p) = \left(\frac{\partial f}{\partial x_1}(p), \frac{\partial f}{\partial x_2}(p), \frac{\partial f}{\partial x_3}(p) \right)^T.$$

It is the zero vector iff p is critical. An *integral line* $z : \mathbb{R} \rightarrow M$ is a maximal path whose velocity vectors agree with the gradient: $\frac{\partial z}{\partial s}(s) = \nabla f(z(s))$ for all $s \in \mathbb{R}$. Each integral

line is open at both ends, and we call $\text{org } z = \lim_{s \rightarrow -\infty} z(s)$ the *origin* and $\text{dest } z = \lim_{s \rightarrow \infty} z(s)$ the *destination* of z . Both are necessarily critical points of f . Integral lines are pairwise disjoint. We consider each critical point as an integral line by itself, and with this stipulation the integral lines partition M . We use them to decompose M into regions of similar flow patterns. The *descending* and *ascending manifolds* of a critical point p are

$$\begin{aligned} D(p) &= \{p\} \cup \{x \in M \mid x \in \text{im } z, \text{dest } z = p\}, \\ A(p) &= \{p\} \cup \{x \in M \mid x \in \text{im } z, \text{org } z = p\}, \end{aligned}$$

where $\text{im } z$ is the image of the path z on M . The descending manifolds of f are the ascending manifolds of $-f$ and, symmetrically, the ascending manifolds of f are the descending manifolds of $-f$. This implies that $D(p) \cap A(p) = p$. It also implies that the two types of manifolds have the same structural properties. Specifically, the descending manifold of a critical point of index i is an open cell of dimension $\dim D(p) = i$. Since the integral lines partition M , so do the descending manifolds. Moreover, they form a complex as the boundary of every cell is the union of lower-dimensional cells that are its *faces*. The ascending manifolds form a dual complex: for critical points p and q of f , $\dim D(p) = 3 - \dim A(p)$, and $D(p)$ is a face of $D(q)$ iff $A(q)$ is a face of $A(p)$.

Morse complexes. A Morse function f is *Morse-Smale* if the descending and ascending manifolds intersect only transversally. Suppose $D(p)$ and $A(q)$ have non-empty common intersection. If $\dim D(p) = 2$ and $\dim A(q) = 1$ then the transversality assumption implies $D(p) \cap A(q) = p = q$. In the more interesting case in which both are 2-manifolds, $A(p)$ and $D(q)$ are faces of $A(q)$ and $D(p)$, and as illustrated in Figure 2, the common intersection is a simple path connecting the two critical points. Following [6],




Figure 2: The common intersection of a descending and an ascending 2-manifold.

we define the cells of the *Morse complex* as the components of the sets $D(p) \cap A(q)$, over all critical points p and q of f . By definition, each cell of the Morse complex is a union of integral lines that all share the same origin q and the same destination p . The dimension of the cell is then the difference between the two indices. We call the cells of dimension 0, 1, 2, and 3 *nodes*, *arcs*, *quadrangles*, and *crystals*. The two-dimensional cells are indeed all quadrangles, but their

boundary can be glued to itself. The prototypical case of a crystal is a cube, which we imagine standing on its tip, but more interesting cases are possible.

3 PL 3-Manifolds

We are interested in algorithms that work for piecewise linear functions obtained from point measurements. In this section, we introduce the necessary terminology, and we discuss some of the difficulties that arise when we transport concepts from the smooth category to the piecewise linear category.

Triangulation. Let K be a simplicial complex that triangulates the 3-manifold M . This means that there is a homeomorphism between M and the underlying space of K , but to simplify the discussion, we assume that M is the underlying space. The complex consists of simplices of dimension 0 through 3, which we refer to as *vertices*, *edges*, *triangles*, and *tetrahedra*. The *star* of a simplex σ consists of all simplices that contain σ as a face, which includes σ itself, and the *link* consists of all faces of simplices in the star that are disjoint from σ :

$$\begin{aligned} \text{St } \sigma &= \{\tau \in K \mid \sigma \leq \tau\}, \\ \text{Lk } \sigma &= \{v \in K \mid v \leq \tau \in \text{St } \sigma, v \cap \sigma = \emptyset\}. \end{aligned}$$

For example, if σ is a vertex then the link is a triangulation of the 2-sphere. Let $f : M \rightarrow \mathbb{R}$ be a continuous function that is linear over every simplex of K . To say this more formally, note that every point x in a simplex is a unique convex combination of its vertices: $x = \sum_{\ell} \lambda_{\ell} u_{\ell}$ with $1 = \sum_{\ell} \lambda_{\ell}$ and $\lambda_{\ell} \geq 0$ for all ℓ . Assuming f is given at the vertices, we then have $f(x) = \sum_{\ell} \lambda_{\ell} f(u_{\ell})$. We will refer to f as a height function and feel free to use relative terms such as ‘higher’ and ‘highest’. It will be convenient to assume that no two vertices have the same height, which can be justified computationally by simulating a perturbation, as described in [5]. The *lower star* (*upper star*) of a vertex u contains all simplices in the star link for which u is the highest (lowest) vertex, and the *lower link* (*upper link*) contains all simplices in the link that are faces of the lower star (upper star):

$$\begin{aligned} \underline{\text{St}} u &= \{\tau \in \text{St } u \mid x \in \tau \implies f(x) \leq f(u)\}, \\ \underline{\text{Lk}} u &= \{v \in \text{Lk } u \mid v \leq \tau \in \underline{\text{St}} u\}, \\ \overline{\text{St}} u &= \{\tau \in \text{St } u \mid x \in \tau \implies f(x) \geq f(u)\}, \\ \overline{\text{Lk}} u &= \{v \in \text{Lk } u \mid v \leq \tau \in \overline{\text{St}} u\}. \end{aligned}$$

By assumption of pairwise different heights, every simplex in K has a unique highest vertex. It follows that the lower stars of the vertices partition K .

Critical vertices. Strictly speaking, critical points for f are not defined, but we may use small bump functions and think of f as the limit of a series of smooth maps. This is the intuition we use to transport concepts and results from

the smooth category to the piecewise linear category. We use lower links and their reduced Betti numbers to distinguish regular from critical vertices and to further classify the latter. The reduced Betti numbers are denoted as $\tilde{\beta}_k$. They are the same as the common un-reduced Betti numbers, except that $\tilde{\beta}_0 = \beta_0 - 1$ for non-empty lower links, and $\tilde{\beta}_{-1} = 1$ for empty lower links [13]. Since lower links are two-dimensional complexes, only $\tilde{\beta}_{-1}$ through $\tilde{\beta}_2$ can be non-zero. As shown in Table 1, the simple critical points are the ones that have exactly one non-zero reduced Betti number, which is equal to one. A vertex that falls outside the

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

Table 1: The classification of regular and simple critical points using reduced Betti numbers.

classification of Table 1 has necessarily $\tilde{\beta}_{-1} = \tilde{\beta}_2 = 0$ and $\tilde{\beta}_0 + \tilde{\beta}_1 \geq 2$. It can be unfolded into $\tilde{\beta}_0$ 1-saddles and $\tilde{\beta}_1$ 2-saddles. One way to do that is to repeatedly cut the link along a circle that intersects the lower link along a single interval. The reduced Betti numbers on the two sides add up to the original ones: $\tilde{\beta}_k = \tilde{\beta}_{kL} + \tilde{\beta}_{kR}$, for $k = 0, 1$. We can always choose the circle such that the sum of the reduced Betti numbers are non-zero on both sides. It then follows that the reduction ends after $\tilde{\beta}_0 + \tilde{\beta}_1 - 1$ cuts and generates $\tilde{\beta}_0$ 1-saddles and $\tilde{\beta}_1$ 2-saddles. We call such a vertex a *multiple saddle*.

Quasi Morse complex. We construct a quasi Morse complex by taking open subcomplexes of K . Intuitively, it is a complex with the same combinatorial structure as the Morse complex. It is a decomposition of space into crystals where the boundary of each crystal is a quadrangulation. The function f has its critical points at the nodes of this complex and is monotonic within all the arcs, quadrangles and crystals. It differs from the Morse complex because the arcs and quadrangles may not be those of maximal ascent and descent.

Let U, V, X , and Y be the set of minima, 1-saddles, 2-saddles, and maxima of a real valued piecewise linear function f defined over a 3-manifold \mathbb{M} . Let P, Q , and R be sets of arcs that connect minima to 1-saddles, 1-saddles to 2-saddles, and 2-saddles to maxima respectively. Let D and A be a set of quadrangles with nodes from U, V, X, V and Y, X, V, X in that order, respectively, around the boundary. We define the *quasi Morse complex* of \mathbb{M} and f as a decomposition of M into cells that satisfies the following properties:

- (i) All nodes are from $U \cup V \cup X \cup Y$, arcs are from $P \cup Q \cup R$, and quadrangles are from $D \cup A$,

- (ii) There are no critical points within the arcs, quadrangles and crystals, and
- (iii) Each arc in Q is on the boundary of four quadrangles, which in a cyclic order alternate between A and D .

Note that the quasi Morse complex can be split into complexes defined by U, D and Y, A . These are the complexes of the descending and ascending manifolds.

Simulating disjointness. Integral lines are not well defined for PL-manifolds. So, following [6], we construct monotonic curves and surfaces that never cross. These curves and surfaces can merge together and fork later. When a curve merges with a surface or two surfaces merge together, we pretend that they remain infinitesimally close to each other without crossing till they either fork or reach a common critical point.

4 Algorithm

In this section, we give an overview of the algorithm and describe some of the fundamental operations, such as computing the link and deciding the criticality of a vertex. Detailed descriptions of how we construct the descending and ascending manifolds will be given in Sections 5 and 6.

Overview. The Morse complex is constructed during two sweeps over the 3-manifold. The first sweep is in the order of decreasing function value, which we refer to as height, and provides the structural information necessary to efficiently compute the descending manifolds. The second sweep is in the order of increasing height, which is the preferred order for computing the ascending manifolds. However, instead of computing the two collections independently, we use the structure provided by the descending manifolds and add the ascending manifolds accordingly.

- Step 1. Construct the complex formed by the descending manifolds.
- Step 2. Compute the intersection curves between the descending and the ascending 2-manifolds.
- Step 3. Construct the ascending manifolds in pieces inside the cells formed by the descending manifolds.

Some routing decisions in Step 1 require the structure of the ascending disk within the star of its origin and so, we compute this structure also. Note that we compute the intersections between the descending and the ascending 2-manifolds before we construct the latter. It is in fact easier to compute these intersections first and then widen them into the ascending 2-manifolds. To streamline our description of the various steps in the algorithm, we use superscripts to represent order information. Specifically, we denote the vertices of K by p^1, p^2, \dots, p^n assuming $f(p^1) > f(p^2) > \dots > f(p^n)$.

Computing links. The link of a vertex p^i needs to be computed at every step during both sweeps over the 3-manifold. A depth first search starting from one of the triangles in $\text{St } p^i$ gives a traversal of all the triangles in the star. Each triangle in the star corresponds to an edge in the link of the vertex and so we also have a traversal of all edges in $\text{Lk } p^i$. This algorithm runs in linear time on the number of triangles in the star (edges in the link). A check to determine whether a vertex in $\text{Lk } p^i$ has a lower (higher) height value compared to p^i can be incorporated into the DFS traversal to compute the lower (upper) link.

Classifying vertices. As discussed in Section 3, the reduced Betti numbers of the lower link can be used to classify the vertices as regular, minimum, 1-saddle, 2-saddle, maximum or a multiple saddle. We now describe the computation of the Betti numbers. The classification and computation of multiplicity for multiple saddles follows immediately.

β_0 can be computed from the lower link by maintaining a union-find data structure on the link vertices during the DFS traversal. Each edge in the link triggers a union operation. β_0 is given by the total number of sets at the end of the traversal. β_2 is non-zero (equals 1) only for a maximum and this case can be easily recognized because $\underline{\text{Lk}} p^i$ and $\text{Lk } p^i$ have the same number of vertices. We compute β_1 from the Euler characteristic χ of $\underline{\text{Lk}} p^i$.

$$\begin{aligned}\chi &= \beta_0 - \beta_1 + \beta_2 \\ &= s_0 - s_1 + s_2\end{aligned}$$

where s_i is the number of i -simplices in $\underline{\text{Lk}} p^i$. We compute the Euler characteristic using the latter expression and use this value to compute β_1 from the former expression.

5 Descending Manifolds

We compute the descending 1- and 2-manifolds simultaneously during one sweep. To simplify the presentation, we will first discuss them separately and restrict our attention to simple critical points. Later in the section, we describe the simultaneous computation.

Descending 1-manifolds. Each descending 1-manifold is an open arc that belongs to a 1-saddle. We denote the arc that belongs to p^i by D_1^i and we call p^i its *root*. As illustrated in Figure 3, the arc descends from its root on both sides and ends at minima of f . (Recall that this is the Morse-Smale condition, descending 1-manifolds can end at other critical points in the non-generic case.) We define the arc as open, so it contains all its vertices except for the minima in which it ends. It is possible that the two branches end at the same minimum, but because they do not contain that minimum, their union is still an open arc and not a closed circle. In the typical (Morse-Smale) case, all vertices of the arc except for its root are regular, but it is also possible that the arc passes




Figure 3: The descending 1-manifold rooted at the 1-saddle p^i . The spheres sketch the links of the root, a regular point, and one of the two minima.

through a 2-saddle or 1-saddle p^j . We have $j > i$ because p^j is necessarily lower than the root, p^i . For the arc it makes little difference whether it passes through a regular or a critical point. However, since p^j starts its own descending manifold, we need to make sure that D_1^i and the descending manifold rooted at p^j are consistent in the sense of *simulated disjointness*. This consistency will be automatic because we extend each arc by adding the edge from the current endpoint to the lowest vertex in its lower link. This choice of extension implies, for example, that once two branches of two arcs merge, they go together until they both end at the same minimum.

We distinguish between the three phases in the construction of the descending 1-manifolds as we sweep the 3-manifold in order of decreasing height: *Starting*, *Expansion*, and *Gluing*. The same three cases also occur in the construction of descending 2-manifolds, and they are processed within the same logical structure. The starting phase applies if the current vertex, p^i , in the sweep is a 1-saddle. In this case, we start the two branches of D_1^i using edges from p^i to the lowest vertex in each component of $\underline{\text{Lk}} p^i$. In the expansion phase, all the descending arcs ending at p^i are extended by adding an edge from p^i to the lowest vertex in $\underline{\text{Lk}} p^i$ unless p^i is a 1-saddle where the edge from p^i to the lowest vertex in one of the components of $\underline{\text{Lk}} p^i$ is added. This choice of lower link component is made such that the descending arc does not intersect the ascending disk originating at the 1-saddle. We describe this choice later in the section. The gluing phase applies if p^i is a minimum in which case we declare p^i to be a node of the Morse complex and glue the descending arcs ending at p^i together. The running time of the algorithm is proportional to the total size of all vertex links plus the total size of the constructed arcs. The former quantity is bounded by the total size of the triangulation, K , but the latter quantity can be more because arcs can share edges.

Structure of descending 2-manifolds. The construction of the descending 2-manifolds is considerably more complicated than that of 1-manifolds. We begin by discussing their structure and by formulating the corresponding invari-

ants to be maintained by the algorithm. Each descending 2-manifold is an open disk that belongs to a 2-saddle. The disk that belongs to p^i is denoted by D_2^i and p^i is its *root*. The disk descends from the root, which is its highest vertex. The boundary of the disk is a circle consisting of descending 1-manifolds that meet at shared minima. The circle might be partially glued to itself along one or more arcs. Note that this is fundamentally different from the case in which the disk folds onto itself. The folding can be simulated away since it does not happen for smooth functions, while the boundary gluing is an inherent feature of descending 2-manifolds. It is important that D_2^i does not contain its boundary, else it would not necessarily be a disk. In the most extreme case, the boundary circle is a single vertex so that its closure is a sphere. This gives the disk the appearance of a pouch.

Beyond being an open disk which descends from its root, we require that the restriction of f to D_2^i has no critical points other than the maximum at p^i . This property is guaranteed by an invariant maintained during the construction. At any moment, we have an open disk whose boundary is partially *final* or *frozen* and partially *unfrozen*. The frozen boundary grows from the empty set to a collection of open segments, which eventually merge to form a complete circle. The unfrozen boundary shrinks from a complete circle to a collection of closed segments, until it eventually disappears.

INVARIANT. Let q be a vertex in the unfrozen portion of the boundary of D_2^i and let qu be an interior edge of the disk. Then u is either an interior vertex or a frozen boundary vertex, and $f(u) > f(q)$.

Note that the Invariant prohibits interior edges that connect two unfrozen boundary vertices. This implies that as long as the entire boundary is unfrozen, there are no interior edges connecting boundary vertices, and all edges descend from the interior to the boundary. Figure 4 illustrates the resulting structure of a descending 2-manifold. A regular vertex u in the restriction of f to the disk is characterized by a non-empty connected lower link. In other words, the edges in the star change between descending from u to descending towards u exactly twice around u . The disk is extended at the




Figure 4: A portion of the triangulation of a partially constructed descending 2-manifold. The arrowheads orient the edges from the higher to the lower endpoints.

highest vertex q that either belongs to the unfrozen boundary or is an endpoint of a frozen boundary segment. In the former case, all interior edges descend towards q . We maintain the Invariant by extending the disk such that all newly added edges descend from q . It follows that the only new interior vertex, which is q itself, is a regular point of f restricted to the disk. In the latter case, we maintain the Invariant by extending the disk such that all newly added edges descend from q .

Starting a disk. Similar to descending 1-manifolds, we start a descending 2-manifold at every 2-saddle, and we extend descending 2-manifolds at all vertices adjacent to unfrozen boundary edges. We begin by discussing how to start a disk. Let $p = p^i$ be a 2-saddle, as shown in Figure 5, and let q be the lowest vertex in its link. By assumption, the lower link forms a belt around the link, and q belongs to that belt. We start the disk D_2^i by constructing a circle in the belt of p , making sure that circle contains q as one of its vertices. Even though we call it a circle, it may fold onto itself, and sometimes such folding is unavoidable.




Figure 5: The disk rooted at p starts by connecting p to a circle in the belt that passes through the lowest vertex q .

There are many ways to construct such a circle. We start with a shortest-path tree from q inside the belt to find a shortest such circle. Assuming all edges have length 1, such a circle minimizes the number of edges. After constructing the tree, we classify non-tree edges in the belt depending on whether or not they separate the two oceans in the link. The circle is then defined by the separating non-tree edge in the belt whose two endpoints minimize the sum of distances to q . We now return to the classification task. The tree cuts the link open but keeps it connected. If we cut along a non-tree edge, we split the link into two disks. If the edge does not separate then one of the disks contains both oceans while the other is contained inside the belt. The latter disk is triangulated, and by construction its triangulation has all vertices on the boundary. We can therefore remove the triangles from the disks by repeated collapsing: at each step remove a triangle that has both edges on the boundary and declare the third edge a new boundary edge. The classification of non-tree belt edges thus proceeds by repeated collapsing, which marks all non-separating edges and leaves all

separating edges unmarked.

The running time of the algorithm is linear in the size of the belt. Indeed, we can construct the shortest-path tree using breadth-first search, collapse triangles using a stack, and select the length minimizing unmarked edge, all in linear time.

Expanding a disk. The interior vertices of a disk are typically regular points of f , although they can also be 1-saddles and 2-saddles. We first consider a regular point $p = p^i$. Since we visit the vertices in the order of decreasing height, p is the highest boundary vertex adjacent to at least one unfrozen boundary edge. Figure 6 illustrates the two possible cases, one in which there are two neighbors a, b connected to p by unfrozen boundary edges, and the other in which there is only one such neighbor, c . The algorithm treats




Figure 6: One expanded disk has p as an unfrozen boundary vertex, and the other has p as an endpoint of a frozen boundary segment.

both cases similarly and simultaneously. Specifically, it constructs a shortest-path tree from the lowest vertex q within the lower link of p . The points a and b belong to the lower link and are therefore vertices of the tree. We connect a to q along the unique path in the tree and extend the corresponding disk by connecting p to the edges of that path. We do the same for b and for all other vertices that are connected to p by unfrozen boundary edges. It is possible that some paths fold onto each other or themselves, and we must keep track of sidedness as before. There is no essential difference in the computations if p is a 2-saddle. We ensure that there are no intersections between descending manifolds by using the same tree for starting and expanding disks into $\underline{\text{Lk}} p^i$. If two disks happen to merge at any stage, they remain so till they reach a minimum. Consider the expansion of any disk into $\underline{\text{Lk}} p^i$. The edges of the newly added triangles form a path within $\underline{\text{Lk}} p^i$ that starts from a leaf or interior vertex of the shortest-path tree and terminates at the root. Naturally, once this path meets a similar path (*i.e.* a path from another disk expanded or started at p^i) they stay merged till the root thereby ensuring that there are no intersections.

The case of a 1-saddle p is more interesting. A disk for which p is an unfrozen boundary vertex whose two neighbors a and b belong to opposite poles attaches to the descending 1-manifold rooted at p . We declare p as frozen for this disk. Finally, we do the same computations as before within each

pole. The case of a newly frozen vertex p appears in both poles, with a single neighbor each.

The running time of the algorithm is linear in the size of the link plus the total size of all extensions. Because of folding, the latter quantity is not necessarily bounded from above by the former.

Simultaneous construction of descending manifolds.

We now describe the construction of descending 1- and 2-manifolds in a single sweep. At each step of the sweep, we process a vertex p^i with a lower link containing β_0 components and β_1 tunnels and perform the following operations in order: start β_1 descending disks, compute ascending disks around the $\beta_0 - 1$ belts in $\overline{\text{Lk}} p^i$, extend all descending disks ending at p^i , start $\beta_0 - 1$ descending arcs, extend all descending arcs ending at p^i , and finally, glue the disks and arcs ending at p^i if it is a minimum.

Starting all descending disks. The procedure described to construct a circle in a belt can be extended to construct circles in each belt within a component of $\underline{\text{Lk}} p^i$ (there are β_1 such belts in $\underline{\text{Lk}} p^i$). After identifying the separating non-tree edge, we mark it and repeat the collapsing process to determine the next separating non-tree edge for a circle around the second belt and so on. Figure 7 shows the circles within the two belts of $\underline{\text{Lk}} p^i$. The descending disks are now given by the triangles in $\text{St } p^i$ that contain the edges in the corresponding circles as faces.




Figure 7: Two disks start from p^i one each around the two belts in $\underline{\text{Lk}} p^i$ and pass through the lowest vertex q .

Starting all ascending disks. In the procedure described for the construction of disks and arcs, we ensured consistency in the sense of simulated disjointness between the descending manifolds. However, we need additional information during the construction in order to prevent non-transversal intersection with the ascending manifolds. Figure 8 illustrates a situation that could lead to inconsistency if not handled properly. The descending 1-manifolds at r should be extended through q_1 in this case, to prevent an intersection. In order to perform this extension, we need the structure of the ascending start disk that originates at p^i .




Figure 8: The structure of the ascending disk starting from p^i , shown here as the shaded disk passing through p^i , is required to decide that the arc along rp^i is to be extended through q_1 and not q_2 .

Ascending start disks are started by constructing circles within the belts of $\overline{\text{Lk}} p^i$ similar to the process of starting descending disks. The circles that we construct now have the additional property that they intersect any descending 2-manifold in $\overline{\text{Lk}} p^i$ at exactly one designated vertex called a *gateway*. This property ensures that there are no non-transversal intersections between ascending and descending 2-manifolds. We modify the circle construction procedure used for starting descending disks in order to construct circles with the above mentioned property. Each descending disk in $\overline{\text{Lk}} p^i$ is considered as a barrier that can be crossed only at its gateway vertex. We simulate this barrier by duplicating all edges and vertices (except the gateway) of the disk in $\overline{\text{Lk}} p^i$ and splitting the set of neighbors of each duplicated vertex between its copies. Figure 9 shows the duplicated vertices and edges of a disk in $\overline{\text{Lk}} p^i$. A gateway can be specified by two adjacent edges in $\overline{\text{Lk}} p^i$ and is chosen for each descending disk in $\overline{\text{Lk}} p^i$ as follows:

1. If both neighbors of p^i along the disk boundary lie in $\underline{\text{Lk}} p^i$ then the gateway is given by the two edges of this disk boundary that are incident on the highest boundary vertex.
2. Check to see if there is a path in $\text{Lk} p^i$ such that
 - a. The end points of the path are in different components of $\underline{\text{Lk}} p^i$ and all other vertices of the path are in $\overline{\text{Lk}} p^i$.
 - b. The path is a union of two or more sub-paths in $\text{Lk} p^i$ that share at most one vertex pairwise and are the intersection curves of a descending disk with $\overline{\text{Lk}} p^i$.

For each such path, designate the two edges in $\overline{\text{Lk}} p^i$ that share the end point of the sub-path with an end point in $\underline{\text{Lk}} p^i$. In Figure 9, d_2, d_3 , and d_4 constitute such a path and we designate (e_1, e_2) as the gateway.

We construct the spanning tree in the graph that simulates the barriers, use the collapsing procedure to classify non-tree belt edges and determine the separating non-tree edge that gives the shortest circle. Note that we only simulate the splitting of vertices and edges in $\overline{\text{Lk}} p^i$ and do not modify the mesh in any manner.

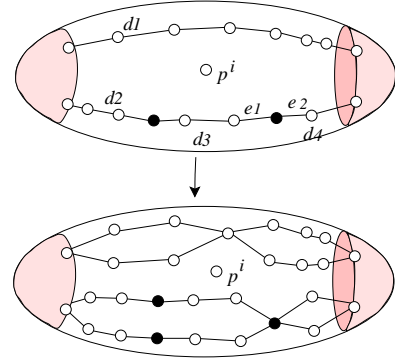


Figure 9: Descending disks d_1, d_2, d_3 , and d_4 in $\overline{\text{Lk}} p^i$ and the simulation of the barrier by splitting vertices and edges except for a gateway. A gateway is chosen for each path between components of $\underline{\text{Lk}} p^i$.

Expanding all descending disks. The procedure that we described previously for expanding disks can be modified to handle a vertex p^i whose lower link has multiple components. If the neighbors of p^i along the boundary of a disk lie within $\underline{\text{Lk}} p^i$ then no modification is necessary. However, if exactly one of the neighbors is frozen (*i.e.* in $\overline{\text{Lk}} p^i$) then we first expand the disk within the component of $\underline{\text{Lk}} p^i$ that contains the other neighbor. Typically this will complete the expansion procedure but it is possible that there are two (or more) disks to be expanded at p^i that share the same neighbor vertex (see Figure 10). In such a case, there is an arc to be extended at p^i into one of the components of $\underline{\text{Lk}} p^i$. We create *spikes* along the arc for all but one of these disks in order to maintain a consistent structure. We can think of the spike as a very thin protrusion at p^i from the boundary of the disk.




Figure 10: There are two disks, d_1 and d_2 , and an arc to be extended at p^i . A spike is added to d_1 and it is extended along with the arc for consistency.

Starting and expanding all descending arcs. We start $\beta_0 - 1$ arcs into the various components of $\underline{\text{Lk}} p^i$ using edges from p^i to the lowest vertex in each component. The ascending disks originating from p^i determine the $\beta_0 - 1$ pairs among the β_0 components of $\underline{\text{Lk}} p^i$. We also use the ascending disks to determine the component of $\underline{\text{Lk}} p^i$ that should be used for extending an arc. Specifically, we choose a component such that the extension of the arc does not cause an

intersection with the ascending disk(s) starting at p^i .

Gluing. If p^i is a minimum then we glue together all the arcs and disks in $\overline{\text{Lk}} p^i$. Two disks merge together first at a 1-saddle on the boundary and are then extended together along the arc on the boundary till a minimum. However, we wait till the minimum before gluing them together. This unifies the process of gluing disks and arcs.

The descending 2-manifolds are split into a set of quadrangles. The arcs along which they are split are exactly the intersection curves between the ascending and descending 2-manifolds. For each disk ending at p^i (a minimum), a new quadrangle is added into the morse complex. Two arcs of this quadrangle are given by descending arcs from the two neighboring 1-saddles of p^i along the boundary of the disk. The remaining two arcs are the intersection curves from the 1-saddles to the 2-saddle that originated the disk. These curves are computed in the next step of the algorithm. The ring of quadrangles around each arc ending at p^i corresponds to the disks that contain the arc in their boundary.

6 Ascending Manifolds

We now describe the construction of the ascending 1- and 2-manifolds. We first build the intersection curves between the ascending and descending 2-manifolds connecting the 1- and 2-saddles. These curves divide the descending disks into quadrangles with nodes at a 2-saddle, a minimum, and two 1-saddles. We then build the ascending arcs that start from 2-saddles and end at maxima. At this stage, the ascending disks are divided into quadrangles with nodes at a 1-saddle, a maximum and two 2-saddles with the bounding arcs already constructed. Finally, we construct the ascending disks by filling up each of these quadrangles. As we did for the descending manifolds, we first discuss the individual steps for handling simple critical points and regular points and later extend to multiple critical points and simultaneous construction. Since the processing follows the same logical structure as in the descending manifold construction, we only give brief descriptions.

Intersection curves. For a smooth Morse-Smale function on a 3-manifold, the intersections of the descending and ascending 2-manifolds are arcs within a descending disk that connect 1-saddles on the boundary of the disk to the origin of the disk, which is a 2-saddle. The invariant that we maintain while constructing descending disks guarantees that the restriction of f to each descending disk has no critical points other than its origin. We start an arc at each 1-saddle in the boundary of a descending disk and follow the edge path within the disk that always connects a vertex to the highest vertex in its link restricted to the disk. These paths may come together at some point, but never cross. Two paths in different disks may also intersect, but this intersection will be resolved when the descending disks get resolved.

Ascending arcs. Ascending 1-manifolds are started, expanded and glued at p^i by a procedure similar to the one used for descending arcs. The only difference is in the expansion phase. We can no longer extend an arc to the highest point in $\overline{\text{Lk}} p^i$ because this may create intersections with descending disks. Instead, we extend the arcs ending at p^i using edges to the highest possible vertex in $\overline{\text{Lk}} p^i$ that does result in an intersection between the arc and a descending disk.

Ascending disks. We construct ascending disks by filling in the quadrangles that they are divided into by the ascending arcs and intersection curves. The concepts of frozen and unfrozen boundary are used again in the expansion of quadrangles and there is a similar invariant property.

Each quadrangle is paired to four critical points: a 1-saddle, two 2-saddles and a maximum. The quadrangle has edge two arcs connecting the origin 1-saddle to two 2-saddles (intersection arcs) and two arcs connecting the two 2-saddles to a maximum (ascending arcs). The quadrangles that are paired to a given 1-saddle will fit together to make its ascending 2-manifold. Quadrangles are constructed by




Figure 11: A portion of the triangulation of a partially constructed quadrangle within an ascending 2-manifold to be expanded at p^i . The arrowheads orient the edges from the higher to the lower endpoints.

a process similar to that of descending disks. During the process, the interior of the part of the boundary that lies in the union of the ascending 1-manifolds and the intersection paths remains frozen, the rest is unfrozen. Figure 11 illustrates the structure of a quadrangle. The notion of invariant is the same, except that the inequality reverses:

INVARIANT. Let q be a vertex in the unfrozen portion of the boundary of a quadrangle of A_2^i and let qu be an interior edge of the quadrangle. Then u is either an interior vertex or a frozen boundary vertex, and $f(u) < f(q)$.

Recall that we already started ascending disks from 1-saddles while computing the descending manifolds. The descending disks divide $\overline{\text{Lk}} p^i$ into regions. The restriction of the ascending disk to each of these regions determines the initial portion of each quadrangle. Expansion proceeds as follows: take the lowest point on the unfrozen part of the boundary of each quadrangle. This point is either regular, a 1-saddle or a 2-saddle. It is also either an interior point of

the boundary, a point of an intersection path, or a point of an ascending 1-manifold. In each case, the adjacent points, a and b , on its boundary either lie in $\overline{\text{Lk}} p^i$ or are frozen. We expand the quadrangle by constructing a shortest path tree rooted at the highest point within the regions of $\overline{\text{Lk}} p^i$ mentioned above and using the paths from a and b (if they are not frozen) to the root. If p^i is a maximum, then we declare it as a node of the morse complex, glue the ascending arcs within each quadrangle ending at p^i and glue together all the quadrangles within ascending disks ending at p^i .

Simultaneous construction of ascending manifolds. We perform all of the above mentioned operations simultaneously during one sweep of the manifold in increasing order of function value. At each step of the sweep, we process a vertex p^i , with an upper link containing β'_0 components, by performing the following operations in order and when applicable: start intersection curves, extend all intersection curves ending at p^i , start $\beta'_0 - 1$ ascending arcs, extend all ascending arcs ending at p^i , identify the quadrangles from all the ascending disks that were started at p^i , extend all quadrangles ending at p^i , and finally, glue the intersection curves, descending disks and arcs ending at p^i if it is a minimum.

All of the above mentioned operations work directly as discussed above even in the simultaneous construction with the only difference being that they are now applied to multiple critical points also.

7 Discussion

This paper introduces Morse complexes for piecewise 3-manifolds as a decomposition of the 3-manifold into crystals with quadrangular faces. It also gives an algorithm to construct a quasi Morse complex for a 3-manifold that guarantees structural correctness. Many interesting issues still remain open.

We can transform the quasi Morse complex into the Morse complex by applying a sequence of operations called handle slides. Figure 12 illustrates this operation within a section of the quasi Morse complex consisting of three simple crystals. As described in [6], using this approach we obtain a Morse complex that is numerically as accurate as the local rerouting operations used to control handle slides. We might have to perform two sets of handles slides, one each on the ascending and descending 1-manifolds. It is unclear how to find and order these two sets of handle slides that bring us closer to the Morse complex.

It is useful to have a hierarchical representation of the Morse complex while working with large data sets. We can create such a hierarchy by performing a sequence of cancellations of pairs of critical points ordered by the persistence algorithm [7]. There are many issues that need to be resolved here: Do all pairs given by the persistence algorithm correspond to edges in the Morse complex? Is it always possible to contract an edge in the Morse complex? How can the

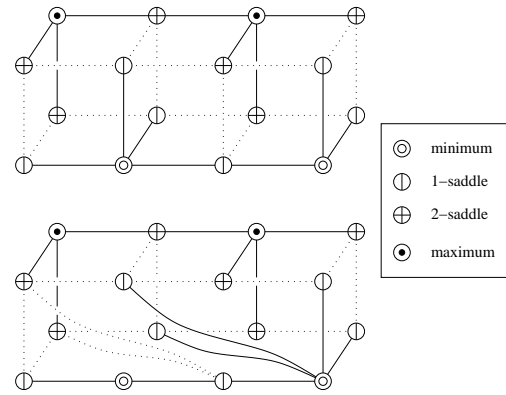


Figure 12: A fraction of the affected Morse complex before and after a handle slide. The arcs between 1-saddles and 2-saddles are dotted, and all other arcs are solid.

Morse function be smoothed so that it corresponds to a cancellation?

References

- [1] H. BLUM *Models for the Perception of Speech and Visual Form*. MIT Press, Cambridge, 1967, 362–380.
- [2] C. L. BAJAJ, V. PASCUCCI AND D. SCHIKORE. Fast isocontouring for improved interactivity. In “Proc. IEEE Sympos. Vol. Viz., 1996”, 39–46.
- [3] H. CARR, J. SNOEYINK AND U. AXEN. Computing contour trees in all dimensions. In “Proc. 11th Ann. SIAM-ACM Sympos. Discrete Algorithms, 2000”, 918–926.
- [4] T. CULVER, J. KEYSER AND D. MANOCHA. Accurate computation of the medial axis of a polyhedron. In “Proc. Sympos. Solid Modeling and Applications, 1999”, 179–190.
- [5] H. EDELSBRUNNER. *Geometry and Topology for Mesh Generation*. Cambridge Univ. Press, England, 2001.
- [6] H. EDELSBRUNNER, J. HARER AND A. ZOMORODIAN. Hierarchical Morse complexes for piecewise linear 2-manifolds. In “Proc. 17th Sympos. Comput. Geom., 2001”, 70–79.
- [7] H. EDELSBRUNNER, D. LETSCHER AND A. ZOMORODIAN. Topological persistence and simplification. *Discrete Comput. Geom.*, to appear.
- [8] A. GUÉZIEC AND R. HUMMEL. Exploiting triangulated surface extraction using tetrahedral decomposition. *IEEE Trans. Visualization Comput. Graphics* **1** (1995), 328–342.
- [9] L. GUIBAS, R. HOLLEMAN AND L. E. KAVRAKI. A probabilistic roadmap planner for flexible objects with a workspace medial axis based sampling approach. In “Proc. IEEE/RSJ Intl. Conf. Intelligent Robots and Systems, 1999”, 254–260.
- [10] W. E. LORENSEN AND H. E. CLINE. Marching cubes: a high resolution 3D surface construction algorithm. *Comput. Graphics* **21**, Proc. SIGGRAPH 1987, 163–169.

- [11] Y. MATSUMOTO. *An Introduction to Morse Theory*. Translated from Japanese by K. Hudson and M. Saito, Amer. Math. Soc., 2002.
- [12] J. MILNOR. *Morse Theory*. Princeton Univ. Press, New Jersey, 1963.
- [13] J. R. MUNKRES. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
- [14] R. L. OGNIWICZ. Skeleton-space: A multiscale shape description combining region and boundary information. In “Proc. Computer Vision and Pattern Recognition, 1994”, 746–751.
- [15] V. PASCUCCI AND K. COLE-MCLAUGHLIN. Efficient computation of the topology of level sets. *Algorithmica*, to appear.
- [16] G. REEB. Sur les points singuliers d’une forme de Pfaff complètement intégrable ou d’une fonction numérique. *Comptes Rendus de L’Académie ses Séances, Paris* **222** (1946), 847–849.
- [17] D. SHEEHY, C. ARMSTRONG AND D. ROBINSON. Shape description by medial axis construction. *IEEE Trans. Visualization Comput. Graphics* **2** (1996), 62–72.
- [18] A. SHEFFER, M. ETZION, A. RAPPOPORT AND M. BERCOVIER. Hexahedral mesh generation using the embedded voronoi graph. *Engineering Comput.* **15** (1999), 248–262.
- [19] D. STORTI, G. TURKIYAH, M. GANTER, C. LIM AND D. STAL. Skeleton-based modeling operations on solids. In “Proc. ACM Sympos. Solid Modeling Applications, 1997”, 141–154.
- [20] S. TARASOV AND M. N. VYALI. Construction of contour trees in 3D in $O(n \log n)$ steps. In “Proc. 14th Ann. Sympos. Comput. Geom., 1998”, 68–75.
- [21] M. VAN KREVELD, R. VAN OOSTRUM, C. BAJAJ, V. PASCUCCI AND D. SCHIKORE. Contour trees and small seed sets for isosurface traversal. In “Proc. 13th Ann. Sympos. Comput. Geom., 1997”, 212–220.
- [22] J. WILHELMS AND A. VAN GELDER. Topological consideration in isosurface generation. *ACM Computer Graphics* **24** (1990), 57–62.
- [23] A. Y. WU, S. K. BHASKAR AND A. ROSENFELD. Computation of geometric properties from the medial axis transform in $O(n \lg n)$ time. *CVGIP* **34** (1986), 76–92.

Appendix A

Table 2 provides a list of notation used in this paper.

$M, f : M \rightarrow \mathbb{R}$	3-manifold, smooth function
$df_p, H(p), \nabla f(p)$	differential, Hessian, gradient
$D(p), A(p)$	descending, ascending manifold
$K, f : M \rightarrow \mathbb{R}$	triangulation, heigh function
$p = p^i, q, r, a, b, q_1, q_2$	vertices in K
$\text{St } p, \underline{\text{St}} p, \overline{\text{St}} p$	star, lower star, and upper star of p
$\text{Lk } p, \underline{\text{Lk}} p, \overline{\text{Lk}} p$	link, lower link, and upper link of p
D_1^i, D_2^j	descending arc, disk
$d1, d2, d3, d4, e1, e2$	edge(s) of descending disk within $\text{Lk } p$

Table 2: Notation for geometric concepts, sets, functions, vectors, variables.