

The Flow Complex: A Data Structure for Geometric Modeling*

Joachim Giesen[†]

Matthias John[‡]

Abstract

Structuring finite sets of points is at the heart of computational geometry. Such point sets arise naturally in many applications. Examples in \mathbb{R}^3 are point sets sampled from the surface of a solid or the locations of atoms in a molecule. A first step in processing these point sets is to organize them in some data structure. Structuring a point set into a simplicial complex like the Delaunay triangulation has turned out to be appropriate for many modeling tasks. Here we introduce the flow complex which is another simplicial complex that can be computed efficiently from a finite set of points. The flow complex turned out to be well suited for surface reconstruction from a finite sample and for some tasks in structural biology. Here we study mathematical and algorithmic properties of the flow complex and show how to exploit it in applications.

1 Introduction

In applications point sets often come unstructured, which does not mean, however, that their distribution is arbitrary. For example the input of the surface reconstruction problem is just a finite set of points, but the points are constrained to lie on the surface of some solid. Another maybe little less intuitive example are cavities in proteins. Atoms of the protein cannot be located in such a cavity. That is, if the point set contains all locations of the atoms in a protein it cannot contain any point that lies in such a cavity. In geometric modeling one asks for a geometric model of such constraints. In the two examples that would either be a model of the surface from which the points are sampled or a model of the cavity. A popular way to model such constraints is as a simplicial complex. In fact surfaces and cavities were successfully modeled as subcomplexes of the Delaunay triangulation of the sample points [1, 2, 3, 5].

In this paper we introduce the flow complex as an alternative to the Delaunay triangulation and show that it is also well suited for surface reconstruction and for the identification of some sorts of cavities in proteins.

The flow complex is closely related to the Delaunay triangulation, but neither complex is a subcomplex of the other. The striking difference is that it seems much easier to extract a surface or cavity model from the flow complex than from the Delaunay triangulation. We will demonstrate this briefly at the end of this paper. In fact, if the point set stems from a surface in \mathbb{R}^3 then the flow complex gives almost a reconstruction.

Though we applied the flow complex successfully in applications the focus of this paper is the theory that lies behind the definition and computation of the flow complex. The starting point of our study is a distance function associated with a finite set of sample points in \mathbb{R}^3 . This function assigns to every point in \mathbb{R}^3 its least distance to any of the sample points. It is intimately related to the Voronoi diagram of the sample points. The distance function has a unique direction of steepest ascent at almost every point. The points where such a direction does not exist are the critical points of the distance function, i.e. its local extrema and saddle points. We study where a point in \mathbb{R}^3 flows if it always follows the direction of steepest ascent of the distance function. It turns out that all points either flow into a local maximum, a saddle point or to infinity. The set of all points that flow into a critical point is called the stable manifold of this critical point. We call the collection of all stable manifolds the flow complex of the sample points.

The main contributions of this paper are new insights in a distance function that was already well studied in the context of Voronoi diagrams and Delaunay triangulations. These insights lead to the definition of a new data structure - the flow complex - that we have successfully applied in surface reconstruction and the identification of pockets in proteins. In particular, we give an efficient algorithm to compute the flow complex and prove its correctness.

2 Diagrams and critical points

In this section we summarize, following [4], the basic notions that we will use throughout the paper.

Voronoi diagram. Let P be a finite set of points in

*Partly supported by the IST Programme of the EU and the Swiss Federal Office for Education and Science as a Shared-cost RTD (FET Open) Project under Contract No IST-2000-26473 (ECG - Effective Computational Geometry for Curves and Surfaces).

[†]Institut für Theoretische Informatik, ETH Zürich, CH-8092 Zürich, Switzerland. Email: giesen@inf.ethz.ch

[‡]Institut für Theoretische Informatik, ETH Zürich, CH-8092 Zürich, Switzerland. Email: john@inf.ethz.ch

\mathbb{R}^3 . The *Voronoi cell* of $p \in P$ is given as

$$V_p = \{x \in \mathbb{R}^3 : \forall q \in P - \{p\}, \|x - p\| \leq \|x - q\|\}.$$

The sets V_p are convex polyhedra or empty since the set of points that have the same distance from two points in P forms a hyperplane. Closed facets shared by two Voronoi cells are called *Voronoi facets*, closed edges shared by three or more Voronoi cells are called *Voronoi edges* and the points shared by four or more Voronoi cells are called *Voronoi vertices*. The term *Voronoi object* can denote either a Voronoi cell, facet, edge or vertex. The *Voronoi diagram* of P is the collection of all Voronoi objects. It defines a cell decomposition of \mathbb{R}^3 .

Delaunay diagram. The *Delaunay diagram* of a set of points P is dual to the Voronoi diagram of P . The convex hull of four or more points in P defines a *Delaunay cell* if the intersection of the corresponding Voronoi cells is not empty and there exists no superset of points in P with the same property. Analogously, the convex hull of three or two points defines a *Delaunay face* or *Delaunay edge*, respectively, if the intersection of their corresponding Voronoi cells is not empty. Every point in P is called *Delaunay vertex*. The term *Delaunay object* can denote either a Delaunay cell, face, edge or vertex. The Delaunay diagram defines a decomposition of the convex hull of all points in B . This decomposition is a triangulation if the points are in general position.

We always refer to the interior and to the boundary of Voronoi-/Delaunay objects with respect to their dimension, e.g. the interior of a Delaunay edge contains all points in this edge besides the endpoints and the interior of a vertex and its boundary are the vertex itself. Furthermore, we always assume general position unless stated differently.

Distance function. Let P be a finite set of points in \mathbb{R}^3 . The distance function induced by P is given as

$$h(x) = \min\{\|x - p\|^2 : p \in P\}.$$

The graph of the distance function h is the lower envelope of a set of paraboloids centered at the points in P . Thus the function h is continuous. It is smooth everywhere besides at points which have the same distance from two or more points, i.e. at points that lie on the boundary of a Voronoi cell.

Regular- and critical points. Following the critical point theory for distance functions developed in Riemannian geometry [7] we define the gradient of a




Figure 1: A one dimensional example that shows the graph of the distance function induced by three points.

distance function h at x as a set $\Gamma(x)$ of unit vectors. In case that h is smooth at x we set $\Gamma(x) = \{\nabla h / \|\nabla h\|\}$ or $\{0\}$ if ∇h vanishes at x . At all other points, i.e. points that lie on the boundary of a Voronoi cell, let $\Gamma(x)$ be the set of unit vectors $\frac{p-x}{\|p-x\|}$ with points $p \in P$ for which $\|x - p\|^2 = h(x)$. Note that in the latter case $\Gamma(x)$ contains more than one vector. The distance function h is regular at x if $\Gamma(x)$ is contained in an open hemisphere of \mathbb{S}^2 . Otherwise we say that it is critical at x . Note that the zero vector is not contained in any open hemisphere of \mathbb{S}^2 . Thus critical points of a smooth function are also critical in this more general setting. We define the index of a critical point as the dimension of the span of the vectors in $\Gamma(x)$. Critical points of index 0 and 3 are local minima and maxima, respectively. Critical points of index 1 and 2 are called saddle points.

We defer the simple proof of the following lemma to the full version of this paper.

LEMMA 2.1. *Let P be a finite set of points such that Voronoi and their dual Delaunay objects intersect in their interiors if they intersect at all. Then the critical points of the distance function h are the intersection points of Voronoi objects V and their dual Delaunay object σ . The index of a critical point is the dimension of σ .*

In the following we always assume that Voronoi and their dual Delaunay objects intersect in their interiors if they intersect at all. Other intersections are degenerate in the sense that they are unstable under small perturbations of the point set.

3 Induced flow

We want to study how the points in \mathbb{R}^3 move if they always follow the direction of steepest ascent of the distance function h . The curve that a point $x \in \mathbb{R}^3$ follows during this motion is called the orbit of x . For smooth distance functions the computation of a single

orbit results in the solution of an ordinary differential equation. Since the distance function h is not smooth everywhere, we cannot apply the theory of ordinary differential equations here. Nevertheless individual orbits can also be computed for h . An essential ingredient of this computation is the notion of the driver of a point $x \in \mathbb{R}^3$. The driver of x is a point $d \in \mathbb{R}^3$ such that the direction of the vector $x - d$ is the direction of the steepest ascent of h at x . Thus knowing the driver of x means knowing in which direction x is going to move.

Driver. Given $x \in \mathbb{R}^3$. Let V be the lowest dimensional Voronoi object in the Voronoi diagram of B that contains x and let σ be the dual Delaunay object of V . The driver of x is the point on σ closest to x .




Figure 2: The critical points of the distance function from Figure 1 and the direction of steepest ascent of the distance function at one point. Note that in one dimension the only critical points of the distance function are local minima \ominus and local maxima \oplus .

The individual orbits of the points in \mathbb{R}^3 can be derived from a so called flow on \mathbb{R}^3 which is a function from $(0, \infty) \times \mathbb{R}^3$ to \mathbb{R}^3 .

Induced flow. The flow ϕ induced by a finite point set P is given as follows: For all critical points x of the height function associated with P we set:

$$\phi(t, x) = x, t \in [0, \infty)$$

Otherwise let y be the driver of x and R be the ray originating at x and shooting in the direction $x - y$. Let z be the first point on R whose driver is different from y . Note that such a z need not exist in \mathbb{R}^3 if x is contained in an unbounded Voronoi object. In this case let z be the point at infinity in the direction of R . We set:

$$\phi(t, x) = x + t \frac{x - y}{\|x - y\|}, t \in [0, \|z - x\|]$$

For $t > \|z - x\|$ the flow is given as follows:

$$\begin{aligned} \phi(t, x) &= \phi(t - \|z - x\| + \|z - x\|, x) \\ &= \phi(t - \|z - x\|, \phi(\|z - x\|, x)) \end{aligned}$$

We will show later that the function ϕ is well defined.

Orbits and fixpoints. Given $x \in \mathbb{R}^3$ and an induced flow ϕ . The curve

$$\phi_x : [0, \infty) \rightarrow \mathbb{R}^3, t \mapsto \phi(t, x)$$

is called the *orbit* of x . A point $x \in \mathbb{R}^3$ is called a *fixpoint* of ϕ if $\phi(t, x) = x$ for all $t \geq 0$.




Figure 3: A two dimensional example that shows four orbits of the flow induced by seven points. The Voronoi diagram of the point set is also shown.

OBSERVATION 3.1. *The following is true.*

- (1) *The fixpoints of ϕ are the critical points of the distance function h .*
- (2) *The orbits of ϕ are piecewise linear curves that are linear in Voronoi objects.*

Because of the first observation we refer to a fixpoint of ϕ as a minimum, saddle or maximum if the corresponding critical point of the distance function is a minimum, saddle point or maximum, respectively.

LEMMA 3.1. *The flow ϕ has no closed orbits.*

Proof. Observe that for every $x \in \mathbb{R}^3$ the distance function associated with P is growing along the orbit ϕ_x for all values $t \in [0, \infty)$ such that $\phi_x(t)$ is not a fixpoint of ϕ . Thus there are no closed orbits possible.

The following lemmas prepare the well-foundedness proof (Theorem 3.1) for the induced flow.

LEMMA 3.2. *All points in the interior of a Voronoi object V have the same driver.*

Proof. Remember that for $x \in \mathbb{R}^3$ its driver is a point $y \in \sigma$ that is closest to x where σ is the Delaunay object dual to the lowest dimensional Voronoi object V that contains x .

Since the affine hulls of V and σ are orthogonal to each other they intersect in exactly one point. Let z be this intersection point. Let $y \in \sigma$ be the point closest to z . Note that it is possible that $y = z$. Since all vectors $y' - z$ with $y' \in \sigma$ are orthogonal to all vectors $x - z$ with x in the affine hull of V we have by Pythagoras theorem for all x in the affine hull of V and all $y' \in \sigma$,

$$\begin{aligned} \|x - y'\|^2 &= \|x - z\|^2 + \|z - y'\|^2 \\ &\geq \|x - z\|^2 + \|z - y\|^2 \\ &= \|x - y\|^2 \end{aligned}$$

Thus y is the closest point on σ for all x in the interior of V , i.e. y is the driver of all these points.

COROLLARY 3.1. *The flow ϕ has only finitely many drivers.*

Power and power distance. We want to assign a weight to every driver and refer to this weight as the power of the driver. Let d be a driver and V be the highest dimensional Voronoi object whose interior points are driven by d . Let σ be the dual Delaunay object of V . Let $x = \operatorname{argmin}_{x' \in V} \|x' - d\|$. The power r_d assigned to d is defined by the following expression,

$$r_d = \|x - d\|^2 - \|x - p\|^2 \text{ for any } p \in P \cap \sigma.$$

Here P is the finite set of points that induces the flow. Observe that

- (1) The definition of the power r_d is independent of the choice of $p \in P \cap \sigma$.
- (2) If $d \in P$ then the power of d is zero.

The power distance of a point $y \in \mathbb{R}^2$ from d is defined as

$$\pi_d(y) = \|y - d\|^2 - r_d.$$

The proof of the following lemma shows that the definition of power for a driver does not depend on the specific choice of x . Every point in V would lead to the same value for the power.

LEMMA 3.3. *Let d be a driver and V be the highest dimensional Voronoi object whose interior points are driven by d . Let σ be the dual Delaunay object of V and let $p \in P \cap \sigma$. The power distance π_d on V is just the squared Euclidean distance from p .*

Proof. Let $y \in V$ arbitrary and $x = \operatorname{argmin}_{x' \in V} \|x' - d\|$. We have by the definition of power,

$$\begin{aligned} \pi_d(y) &= \|y - d\|^2 - r_d \\ &= \|y - d\|^2 - \|x - d\|^2 + \|x - p\|^2 \end{aligned}$$

Remember that the driver d is the closest point in σ to the intersection point z of the affine hulls of V and σ . The vectors $d - z$ and $x - z$ are orthogonal to each other, because we have $d \in \sigma$ and $x \in V$. Hence we get

$$\|x - d\|^2 = \|x - z\|^2 + \|z - d\|^2,$$

which follows from Pythagoras' theorem. From an analogous reasoning we obtain

$$\|x - p\|^2 = \|x - z\|^2 + \|z - p\|^2.$$

This leads to

$$\pi_d(y) = \|y - d\|^2 - \|z - d\|^2 + \|z - p\|^2$$

We apply an analogous reasoning a third and a fourth time to get

$$\|y - d\|^2 = \|y - z\|^2 + \|z - d\|^2$$

and

$$\|y - p\|^2 = \|y - z\|^2 + \|z - p\|^2$$

which leads to

$$\pi_d(y) = \|y - p\|^2 = \pi_i(y).$$

LEMMA 3.4. *The power of the drivers in the sequence of drivers associated with an orbit ϕ_x is always monotone decreasing, i.e. every driver occurs at most once in the sequence of drivers associated with ϕ_x .*

Proof. For every $z \in \mathbb{R}^3$ let D_z be the set of potential drivers of z , i.e. all drivers contained in the Delaunay object dual to the lowest dimensional Voronoi object V that contains z .

Let us collect some properties of D_z .

- (1) According to Lemma 3.2 the set D_z is always finite.
- (2) For every z' sufficiently close to z we have $D_{z'} \subset D_z$, because V has the smallest dimension among all Voronoi objects that have a non empty intersection with a sufficiently small neighborhood of z and all these Voronoi objects have a non empty intersection with V , i.e. the dual Delaunay objects of all these Voronoi objects are contained in the dual Delaunay object of V .

Let $t \in (0, \infty)$ be such that the driver d of $\phi_x(t)$ is different from the driver of $\phi_x(t - \varepsilon)$ for all sufficiently small $\varepsilon > 0$. The orbit ϕ_x is continuous by definition. Thus we have for a fixed $\varepsilon > 0$ which is sufficiently small that the driver d' of $\phi_x(t - \varepsilon)$ is a potential driver of $\phi_x(t)$. That is,

$$\|\phi_x(t) - d\| < \|\phi_x(t) - d'\|.$$

On the other hand we have from Lemma 3.3

$$\pi_d(\phi_x(t)) = \pi_{d'}(\phi_x(t))$$

since $\phi_x(t)$ lies in the boundary of the lowest dimensional Voronoi object that contains $\phi_x(t - \varepsilon)$. This and the definition of power distance imply that the power of d has to be smaller than the power of d' .

THEOREM 3.1. *The mapping ϕ is a well defined function, i.e. it is defined for every $(t, x) \in [0, \infty) \times \mathbb{R}^3$.*

Proof. By construction there exists an $\varepsilon > 0$ for every $x \in \mathbb{R}^3$ such that ϕ is defined on $\{x\} \times [0, \varepsilon)$, i.e. ϕ_x is defined on $[0, \varepsilon)$. This already implies that ϕ_x is defined on the whole of $[0, \infty)$ since ϕ_x is locally defined for every driver and the sequence of drivers associated with ϕ_x is finite according to Corollary 3.1 and Lemma 3.4.

4 Stable manifolds and the flow complex

We are not really interested in the individual orbits of the points, but want to look on the flow on a coarser level. Therefore we group all points together that flow into the same fixpoint of the flow.

Stable manifolds. Given an induced flow ϕ . The stable manifold $S(x)$ of a fixpoint $x \in \mathbb{R}^3$ contains all points that flow into x , i.e.

$$S(x) = \{y \in \mathbb{R}^3 : \lim_{t \rightarrow \infty} \phi_y(t) = x\}.$$

Instead of directly working with stable manifolds of critical points we introduce a smoothed version which has nicer properties. Smoothing essentially means taking the closure of the stable manifold. It will become obvious later why the following definition is more complicated.

Smoothed stable manifolds. Let x be a critical point of index i of an induced flow. Let S be the set of points in $S(x)$ that have a neighborhood in $S(x)$ that is homeomorphic to an open subset of \mathbb{R}^d , $d = i+1, \dots, 3$. Let S' be the boundary of $S(x) - S$ in \mathbb{R}^3 . The smoothed stable manifold of x is the set $S^*(x) = (S(x) - S) \cup S'$.




Figure 4: A two dimensional example of a stable manifold of a maximum \oplus (on the left) and its smoothed version on the right.

An induced flow in \mathbb{R}^3 has four different types of fixpoints, local minima, saddle points of index 1, saddle points of index 2 and local maxima. In the following we are going to characterize the smoothed stable manifolds of the four different types of fixpoints.

OBSERVATION 4.1. *The smoothed stable manifold of a local minimum m contains just the point m itself since no other point flows into m .*

It turns out that the stable manifold of an index 1 saddle point is always a Gabriel edge and vice versa.

Gabriel graph. The Gabriel graph of a finite set of points P in \mathbb{R}^3 is given as follows: Its vertices are the points in P and its edges are given by Delaunay edges that intersect their dual Voronoi facet. The edges of the Gabriel graph are called Gabriel edges. The Gabriel graph is always connected, because it contains the minimum spanning tree of P .

LEMMA 4.1. *Let s be an index 1 saddle of ϕ . The smoothed stable manifold $S^*(s)$ of s is a Gabriel edge and every Gabriel edge is the smoothed stable manifold of some index 1 saddle.*

Proof. Every Delaunay edge that contains some index 1 saddle is a Gabriel edge and vice versa by Lemma 2.1 and the definition of Gabriel edges, respectively. It remains to show that the smoothed stable manifold $S^*(s)$ of every index 1 saddle s is a Gabriel edge.

Let E be the Delaunay edge that contains s . By the definition of ϕ all interior points of E belong to $S(s)$. Hence $E \subset S^*(s)$. We want to show that $E = S^*(s)$. Assume there exists $x \in S^*(s)$ with $x \notin E$. By the definitions of $S^*(s)$ and T that means that there exists $y \in S(s)$ but $y \notin E$.

Since E contains s the orbit of y cannot be disjoint from the interior of E , i.e. ϕ_y has to join E at some point z for the first time and to stay in E afterwards. The point y must be driven into z by one of the potential drivers of z . These are all drivers that are contained in the Delaunay object dual to the lowest dimensional Voronoi object that contains z . The (potential) driver of the interior points in $E - \{s\}$ is one of the endpoints of E . The point s has both of these endpoints and s itself as potential drivers. Thus y must be driven into z by one of the endpoints of E , because s itself cannot drive y into s . But that means that z cannot be the first point on the orbit ϕ_y that is contained in E . That is a contradiction. Thus we have $S^*(s) = E$.

LEMMA 4.2. *Let s be an index 2 saddle of ϕ . If the stable manifold $S(s)$ of s does not contain a Voronoi vertex then $S^*(s)$ is a piecewise linear surface with boundary. The boundary of the surface is made up from Gabriel edges.*

Proof. Assume that s is a saddle of index 2 such that $S(s)$ does not contain a Voronoi vertex. According to Lemma 2.1 we have that s is the intersection point of a Delaunay facet F and its dual Voronoi edge E .

We are going to construct the surface explicitly. We start with the construction of a polygon P whose interior points flow into s . This polygon contains s and is contained itself in F . Since we assume general position there are three Voronoi facets incident to E . The drivers of these facets are points on Delaunay edges in the boundary of F . Such a driver might be a saddle of index 1. Consider a driver d which is not a saddle of index 1. The line segment that connects d with s is contained in F and intersects the boundary of the corresponding Voronoi facet in two points, namely in s and in a second point s' . We get a polyline, i.e. a simple piecewise linear curve, from the two segments that connect s' to the two Delaunay vertices incident to the Delaunay edge that contains d . If the driver of the Voronoi facet is a saddle of index 1 we take its dual Delaunay edge as the polyline. That is, we get three polylines all contained in F , one for each Voronoi facet incident to E . Let P be the polygon whose boundary is given by these polylines. P is contained in F and all its interior points flow into s , see Figure 5. It can be triangulated by connecting s with the points s' and the Delaunay vertices incident to F .

Let s' be a point as constructed above for a Voronoi facet that is not driven by a saddle of index 1. By construction s' is contained in a Voronoi edge E' . Furthermore, by our assumption it has to be an interior point of E' . Since we assume general position E' is incident to three Voronoi facets. For one of these Voronoi facets we have




Figure 5: Two examples of polygons that are contained in F and whose interior points flow into s . The polygon on the left has one index 1 saddle point on its boundary.

already computed a polyline. For the remaining two we do it exactly the same way we did it above for P . Thus we have again three polylines, one for each Voronoi facet incident to E' . Two of these polylines always intersect in their common Delaunay vertex. That is, the three polylines together form a polyline which is homeomorphic to S^1 . The latter polyline need not be contained in a hyperplane but it can be triangulated by connecting the point s' with newly computed points s'' and to the Delaunay vertices incident to the Delaunay facet dual to E' . This gives us a new triangulated surface patch whose interior points all flow into s via s' .

We continue with the above construction until there are no more points s' left for which we have not already constructed a surface patch. The constructed surface patches cannot intersect each other or themselves, because this would mean that there are points which flow into two directions. This is impossible by the definition of ϕ . Hence the construction gives us a triangulated surface T with boundary whose points all belong to $S^*(s)$ though not all of them belong to $S(s)$. By construction the boundary of the surface is made up from Gabriel edges, i.e. Delaunay edges that contain an index 1 saddle.

So far we know that $T \subset S^*(s)$. Next we want to show $T = S^*(s)$. Assume there exists $x \in S^*(s)$ with $x \notin T$. By the definitions of $S^*(s)$ and T this implies that there exists $y \in S(s)$ but $y \notin T$.

Since $s \in T$ the orbit ϕ_y of y cannot be disjoint from the interior of the surface T , i.e. ϕ_y has to join T at some point z and to stay in T afterwards. The point y must be driven into z by one of the potential drivers of z . These are all drivers that are contained in the Delaunay object dual to the lowest dimensional Voronoi object that contains z . The point z cannot be an interior point of a Voronoi cell, because two orbits cannot meet in the interior of a Voronoi cell by the definition of ϕ . That is, the point z is either an interior- or a boundary point of a Voronoi facet, i.e. an interior point of a Voronoi edge.

Let us first discuss the case that z is an interior point of a Voronoi facet. The potential drivers of z are the driver of the Voronoi facet that contains z and the endpoints of the Delaunay edge dual to this Voronoi facet. But by construction T contains the intersection of a small neighborhood of z with the line segments connecting z to these three drivers. Hence z cannot be the first point on the orbit ϕ_y which is contained in T . That is a contradiction. Now assume that z is an interior point of a Voronoi edge. The only interior points of Voronoi edges contained in T are the points s' in the above construction. The Delaunay facet dual to the Voronoi edge that contains z contains six drivers. These drivers are its three incident Delaunay vertices and one driver on every of its three incident Delaunay edges. It does not contain a driver in its interior since by construction it is not intersected by the affine hull of its dual Voronoi edge. This implies also that the driver that we used to construct s' cannot drive y into z . This driver is one of the drivers on the Delaunay edges. But by construction T contains the intersection of a small neighborhood of z with the line segments connecting z to the remaining four drivers. Hence z cannot be the first point on the orbit ϕ_y which is contained in T . That is a contradiction. Hence $S^*(s) = T$.

Note that we do not claim that the surface is homeomorphic to a disk. In fact, the surface need not be simply connected, i.e. it can have holes.

The example in Figure 6 shows that the lemma does no longer hold if $S(s)$ contains a Voronoi vertex. But the latter is a very degenerate situation, i.e. it is not stable under small perturbations of the initial point set. This example also motivates our definition of smoothed stable manifolds which basically says that we only keep the boundaries of three dimensional inflow regions of an index 2 saddle point.




Figure 6: An example of an index 2 saddle point \odot whose stable manifold contains a Voronoi vertex \circ . The inflow region of this Voronoi vertex is three dimensional. That is, the shown stable manifold contains two and three dimensional parts.

In the following we assume that none of the stable manifolds of index 2 saddle points contains a Voronoi vertex.

Flow complex. Given a finite set of points in \mathbb{R}^3 . We call the simplicial complex build by the Gabriel graph and the triangulated surfaces from Lemma 4.2 the flow complex of the point set.

It remains to characterize the stable manifolds of the maxima. There we need the following lemma whose technical proof we defer to the full version of this paper.

LEMMA 4.3. *Let K be the set of all points that do not flow to infinity, i.e.*

$$K = \mathbb{R}^3 - \{x \in \mathbb{R}^3 : \forall n \in \mathbb{N} \exists t_n \geq 0 \text{ s.t.} \\ \forall t > t_n \|\phi_x(t) - 0\| > n\}.$$

Let S bet the set of all saddle points, M be the set of all maxima and M' be the set of all minima of ϕ . The set $C = \{S(y) : y \in S \cup M\}$ covers $K - M'$, i.e. every $x \in K$ that is not a minimum is contained in exactly one stable manifold $S(y)$ of some saddle point or maximum y of ϕ .

THEOREM 4.1. *The smoothed stable manifolds of the maxima of ϕ are exactly the closures of the bounded regions of the stable flow complex provided no stable manifold of an index 2 saddle contains a Voronoi vertex.*

Proof. The Euler characteristic of the flow complex can be computed as an alternating sum of Betti numbers $\chi = -\beta_3 + \beta_2 - \beta_1 + \beta_0$, see [8]. All Betti numbers are non negative. The third Betti number β_3 is zero for trivial reasons. The second Betti number β_2 counts the number of non bounding shells, i.e. the number of bounded regions of the flow complex. The zeroth Betti number counts the number of connected components. The connectedness of the Gabriel graph and the construction of the surfaces in Lemma 4.2 imply that the stable flow complex is connected. Hence $\beta_0 = 1$. Combining all these informations we get

$$\chi - 1 \leq \# \text{ bounded regions of the flow complex.}$$

There is another way to compute the Euler characteristic of a simplicial complex. That is $\chi = \sum_{i=0}^3 (-1)^i c_i$, where c_i denotes the number of i -dimensional simplices in the complex. By construction the flow complex is a simplicial complex and we have $c_3 = 0$. To compute the remaining numbers let us recall the construction of the surfaces in Lemma 4.2. We will count how many vertices, edges and triangles besides vertices and edges on its boundary every surface contributes to the computation of the Euler characteristic χ . Remember that we

constructed these surfaces from patches. We first constructed a polygon P with one inner vertex (the saddle of index 2) and as many inner edges as triangles. We will later take care of the vertices and edges in the boundary of P . Thus the contribution to χ is 1. Every other patch has one vertex and two edges incident to an already constructed patch plus one more triangle than inner edges. Hence the contribution of any such patch to χ is zero. The only vertices and edges of the flow complex that we have not considered so far are exactly the original points and the Gabriel edges. The original points are just the minima of the flow ϕ and the Gabriel edges are in a one-to-one correspondence with the saddles of index 1. The reasoning above shows that all other contributions to χ can be subsumed in a contribution of 1 for every saddle of index 2. Thus we have

$$\begin{aligned} \chi &= \# \text{ saddles of index 2} \\ &\quad - \# \text{ saddles of index 1} \\ &\quad + \# \text{ minima.} \end{aligned}$$

From a generalization of a theorem of Siersma [9] we know that

$$\begin{aligned} -1 &= \# \text{ maxima} - \# \text{ saddles of index 2} + \\ &\quad \# \text{ saddles of index 1} - \# \text{ minima} \end{aligned}$$

Combining all these equations leads to

$$\begin{aligned} \# \text{ maxima} \\ \leq \# \text{ bounded regions of the flow complex} \end{aligned}$$

To establish equality in the last inequality it remains to show that every bounded region of the flow complex contains a maximum. Let x be an interior point of such a bounded region K . The point x cannot flow to infinity, because to do so its flow has to hit the flow complex which means that it belongs to this complex and flows into a saddle. We know from Lemma 4.3 that every non critical point that does not flow to infinity is contained in the stable manifold of either a saddle or a maximum. We know that x cannot flow into a saddle, because otherwise it would be contained in the flow complex. Thus x must flow into a maximum in the interior of K . Hence every bounded region of the flow complex contains at least one maximum.

We conclude that the closures of the bounded regions of the flow complex are exactly the smoothed stable manifolds of the maxima of ϕ .

Especially, this theorem states that that boundary of the smoothed stable manifold of a maximum is made up from smoothed stable manifolds of index 2 saddle points. That is, the smoothed stable manifolds of critical points of different index have a nice recursive structure.

5 Algorithm

The proof of Theorem 4.2 leads immediately to an efficient algorithm that computes the flow complex of a finite set P of points in \mathbb{R}^3 . In the following pseudo code we show how to compute the triangles of the flow complex which is the most complicated part in the computation of this complex. The purpose of this pseudo code is also to present the construction in the proof of Theorem 4.2 in a more formal and compact way.

STABLEFLOWCOMPLEX(P)

```

1   $F := \emptyset$ 
2  compute Voronoi- and Delaunay diagram of  $P$ .
3  compute the set  $S$  of index 2 saddles.
4  for each  $s \in S$  do
5     $f :=$  Delaunay facet that contains  $s$ .
6     $Q := \emptyset$ 
7    for each Delaunay edge  $e$  incident to  $f$  do
8       $Q.\text{push}((s, e))$ 
9    end for
10   while  $Q \neq \emptyset$ 
11      $(v, e) := Q.\text{pop}$ 
12      $u, w :=$  endpoints of  $e$ .
13     if  $e$  contains a saddle of index 1 do
14        $F := F \cup \{uvw\}$ 
15     else do
16        $f :=$  Voronoi facet dual to  $e$ .
17        $d :=$  driver of the interior of  $f$ .
18        $v' :=$  first point on the ray from  $d$  to  $v$ 
19         that is contained in  $f$ .
20        $F := F \cup \{vv'u, vv'w\}$ 
21        $f' :=$  Delaunay facet dual to the
22         Voronoi edge that contains  $v'$ .
23       for each edge  $e' \neq e$  incident to  $f'$  do
24          $Q.\text{push}((v', e'))$ 
25       end for
26     end while
27   end for
28 return  $F$ 

```

6 Applications

6.1 Surface reconstruction Surface reconstruction is a powerful modeling paradigm. To create a model of some solid in \mathbb{R}^3 one can just sample its surface and apply a surface reconstruction algorithm to the sample. The most common model produced in surface reconstruction is a piecewise linear manifold. In [6] we show how to extract a two dimensional manifold from the flow complex. The manifold extraction turns out to be much simpler as for Delaunay based methods and works very well in practice.



Figure 7: This example shows the reconstruction of the HAPPY BUDDHA data set from the Stanford 3D scanning repository. The flow complex is shown on the left and the manifold reconstruction on the right.

6.2 Pockets in Proteins A union of balls in \mathbb{R}^3 where each ball represents an atom is a natural way to model proteins. Such models are called space filling models. A ball is represented by a pair $(z, r) \in \mathbb{R}^3 \times [0, \infty)$, where $z \in \mathbb{R}^3$ denotes the center of the ball and r denotes its radius. In space filling models the balls are centered at the locations of the corresponding atoms. One can get these locations for example from X-ray diffraction of the crystallized protein. The radius of the ball is usually taken to be the van der Waals radius of the corresponding atom plus the radius of some solvent molecule also modeled as a ball.

The concept of pockets was introduced in [5] to model essential cavities in proteins. The definitions given in [5] are all based on the (weighted) Delaunay triangulation of a set of balls, but inspired our work on the flow complex. Let B be the set of balls used in the space filling model of some protein. A *void* is defined as a compact connected region in the complement of the union of balls in B . If we let the balls in B grow new voids are created and existing ones get destroyed. Finally all voids get destroyed. See Figure 9 for two snapshots of such a growth process for a union of disks in the plane.

While the balls grow a void shrinks until it finally consists of only a single point before it vanishes. We can define a distance function from a finite set of balls similarly to the definition of the distance function from a finite set of points that we have considered so far. The point where a void vanishes is a maximum of the distance function associated with B . We group such




Figure 9: On the left, a space filling model of a molecule made up from six atoms in two dimensions. The atoms of this molecule do not define a void. On the right, if we grow the disks a void emerges which gets destroyed if we grow the disks further. The point at x is a maximum, i.e. at x another void which was created earlier has already been destroyed.

maxima together if they can be connected by a path in the complement of the space filling model inside the convex hull of the ball centers, i.e. in the complement of the space occupied by the balls in B which is contained in the convex hull of the ball centers. Essentially, these groups of maxima define a pocket. There are different approaches to assign a shape to a pocket. The approach we take here is based on extension of the flow complex for a finite set of balls. The shape we assign to a pocket is not essentially different from the shape assigned in [5], but we think that in the context of flow diagrams the concept of pockets is easier to grasp.

We partition the set of critical points of the flow induced by B into two sets. The first set contains all critical points that are covered by balls from the space filling model, i.e. every such critical point is contained in at least one ball from B . The critical points in the second set are all critical points not contained in the first set. We call the critical points from the first set *negative* and the points from the second set *positive*. The maxima at which a void vanishes are exactly the positive maxima. We are going to use the stable manifolds of the positive critical points to define a pocket. The collection of all these stable manifolds is called *restricted flow complex*. See Figure 10 for an example.

Pocket. A pocket is a maximal connected component of the restricted flow complex, i.e. the component is connected and there is no larger connected component in the restricted flow complex that contains it.




Figure 8: On the left, a pocket in a fatty acid that does not have an opening to the outside. In the middle a pocket with one opening to the outside in some receptor protein. On the right, a pocket with two openings to the outside in a Gramicidin protein. Here the boundaries of the pockets are visualized which were determined using the characterization in Theorem 4.1.




Figure 10: An example of a restricted flow complex. This complex contains the stable manifolds of two saddles \ominus and two maxima \oplus . It defines one pocket. At both maxima voids get destroyed if we grow the disks.

7 Conclusion and future work

We have introduced a new family of simplicial complexes called flow complex. Their definition was inspired by the work done in [5] and thus it is not surprising that the flow complex can be used to model pockets in proteins. But we think that the concept of pockets is easier to grasp in the context introduced here. More surprisingly the flow complex turned out to be well suited for surface reconstruction.

We have implemented the algorithm FLOWCOMPLEX for weighted points. This algorithm is conceptually similar to the algorithm presented here, but more involved in parts. The algorithm performs well in practice, e.g. the computation of flow complex of 144,647 points sampled from the surface of the BUDDHA model takes 69 seconds on a 480 Mhz Sun Ultra Sparc II workstation. At the moment we know neither the combinatorial nor the algorithmic complexity of the flow diagram

though there is a trivial $O(n^6)$ upper bound. Our implementation suggests that this bound is far too large. On the other hand we know that in two dimensions the complexity of the flow diagram is larger than that of the Delaunay triangulation, namely $\Theta(n^2)$.

References

- [1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.*, **22**, pp. 481–504, (1999)
- [2] N. Amenta, S. Choi, T.K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th. ACM Sympos. Comput. Geom.*, pp. 213–222, (2000)
- [3] T.K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. In *Proc. 17th. ACM Sympos. Comput. Geom.*, pp. 257–263, (2001)
- [4] H. Edelsbrunner. Geometry and Topology for Mesh Generation. *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press (2001)
- [5] H. Edelsbrunner, M. A. Facello and J. Liang. On the definition and the construction of pockets in macromolecules. *Discrete Apl. Math.* **88**, pp. 83–102, (1998)
- [6] J. Giesen and M. John. Surface reconstruction based on a dynamical system. To appear in *Proc. Eurographics 2002*
- [7] K. Grove. Critical Point Theory for Distance Functions. In *Proceedings of Symposia in Pure Mathematics* **54**(3), pp. 357–385, (1993)
- [8] J.R. Munkres. Elements of Algebraic Topology. Perseus Book Publishing, L.L.C. (1984)
- [9] D. Siersma. Voronoi Diagrams and Morse Theory of the Distance Function. In *Geometry in Present Day Science*, O.E. Barndorff-Nielsen and E.B.V. Jensen (eds.), pp. 187–208, World Scientific (1999)