Extreme Elevation on a 2-Manifold *

Pankaj K. Agarwal[†], Herbert Edelsbrunner[‡], John Harer[§] and Yusu Wang[†]

Abstract

Given a smoothly embedded 2-manifold in $\mathbb{R}^3$, we define the elevation of a point as the height difference to a canonically defined second point on the same manifold. Our definition is invariant under rigid motions and can be used to define features such as lines of discontinuous or continuous but non-smooth elevation. We give an algorithm for finding points of locally maximum elevation, which we suggest mark cavities and protrusions and are useful in matching shapes as for example in protein docking.

1 Introduction

The starting point of our work is the desire to identify features that are useful in finding a fit between solid shapes in $\mathbb{R}^3$. We are looking for cavities and protrusions and a way to measure their size. The problem is made difficult by the interaction of these features, which typically exist at various scales. We therefore take an indirect approach, defining a real-valued function on the surface that is sensitive to the features of the shape. We call this the elevation function because it has similarities to the elevation measured on the surface of the Earth, but the problem for general surfaces is more involved and the analogy is not perfect.

Related work in protein docking. The primary motivation for work reported in this paper is protein docking, which is the computational approach to predicting protein interaction, a biophysical phenomenon at the very core of life. The phenomenon is clearly important, and the interest in protein docking is correspondingly wide-spread. We refer to survey articles by Elcock et al. [10], Halperin et al. [14], and Janin

and Wodak [16]. The idea of docking by matching cavities with protrusions goes back to Crick [6] and Connolly [5]. Connolly also introduced the idea of using the critical points of a real-valued function defined on the protein surface to identify cavities and protrusions. The particular function he used is the fraction of a fixed-size sphere that is buried inside the protein volume as we move the sphere center on the protein surface. In the limit, when the size of the sphere goes to zero, this function has the same critical points as the mean curvature function [3]. A similar but different function suggested for the same purpose is the atomic density [18]. Here we take the buried fraction of the ball bounded by the sphere but we also vary its radius from zero to about ten Angstrom. At every point of the protein surface, the function value is the fraction of buried volume averaged over the different size balls centered at that point.

Results. The main contribution of this paper is the description and computation of a new type of feature points that mark extreme cavities and protrusions on a surface embedded in $\mathbb{R}^3$. More specifically,

- we extend the concept of topological persistence [8] to form a pairing between all critical points of a function on a 2-manifold embedded in $\mathbb{R}^3$;

- we use the pairings obtained for a 2-parameter family of height functions to define the elevation function on the 2-manifold;

- we classify the generic local maxima of the elevation function into four types;

- we develop and implement an algorithm that computes all local maxima of the elevation function defined on a triangulated surface in $\mathbb{R}^3$.

The elevation differs from Connolly's and the atomic density functions in two major ways: it is independent of scale, and it provides, beyond location, estimates for the direction and size of shape features. Both additional pieces of information are useful in shape characterization and matching. Examples of the four generic types of local maxima are illustrated in Figure 1. In each but the first case, the maximum is obtained
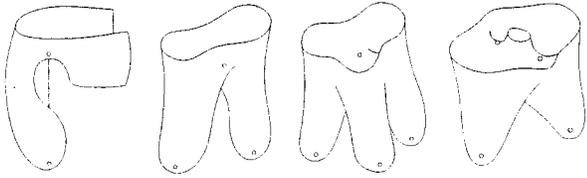
Figure 1: From left to right: a one-, two-, three-, and four-legged local maximum of the elevation function. In the examples shown, the outer normals at the endpoints of the legs are all parallel (the same). Each of the four types also exists with anti-parallel outer normals.

at an ambiguity in the pairing of critical points. In all cases, the endpoints of the legs share the same normal line, and the legs have the same length if measured along that line. The case analysis is delicate and aided by a transformation of the original 2-manifold to its pedal surface, which maps tangent planes to points and represents points with common tangent planes as self-intersections. The algorithm we describe for enumerating all local maxima is inspired by our analysis of the smooth case but works on piecewise linear data.

Outline. Section 2 defines the pairing of the critical points. Section 3 introduces the height and elevation as functions on a 2-manifold. Section 4 describes a dual view of these concepts. Section 5 uses surgery to make elevation continuous. Section 6 characterizes the four types of generic local maxima. Section 7 sketches an algorithm for enumerating all local maxima. Section 8 concludes the paper.

2   Pairing

The elevation function is based on a canonical pairing of the critical points, which we describe in this section.

Traditional persistence. Let $\mathbb{M}$ be a compact, connected, and orientable 2-manifold and $f : \mathbb{M} \to \mathbb{R}$ a smooth function. A point $x \in \mathbb{M}$ is *critical* if the derivative of f at x is identically zero, and it is *non-degenerate* if the Hessian at the point is invertible. It is convenient to assume that f is generic:

    I. all critical points are non-degenerate;

    II. all critical points have distinct function values.

A function that satisfies Conditions I and II is usually referred to as a *Morse function* [17]. It has three types of critical points: *minima*, *saddles* and *maxima*, distinguished by the number of negative eigenvalues of the Hessian. Imagine we sweep $\mathbb{M}$ in the direction of increasing function value, advancing across a level set of closed curves. We write $\mathbb{M}_a = \{x \in \mathbb{M} \mid f(x) \leq a\}$ for the swept portion of the 2-manifold. This portion changes the topology whenever the level set passes through a critical point. A component of $\mathbb{M}_a$

starts at a minimum and ends when it merges with another, older component at a saddle. A hole in the 2-manifold starts at a saddle and ends when it is closed off at a maximum. After observing that each saddle either merges two components or starts an new hole, but not both, it is natural to pair up the critical point that starts a component or a hole with the critical point that ends it. In a nutshell, this is the idea of topological persistence introduced in [8]. It is clear that a small perturbation of the function that preserves the sequence of critical events does not affect the pairing, other than by perturbing each pair locally. The method pairs all critical points except for the first minimum, the last maximum, and the $2g$ saddles starting the $2g$ cycles that remain when the sweep is complete, where g is the *genus* of $\mathbb{M}$. These $2 + 2g$ unpaired critical points are the reason we need an extension to the method, which we describe next.

Extended persistence. It is natural to pair the remaining minimum with the remaining maximum. The remaining $2g$ saddles are paired in a way that reflects how they introduce cycles during the sweep. This pairing is best described using the *Reeb graph* obtained by mapping each component of each level set to a point, as illustrated in Figure 2. As proved
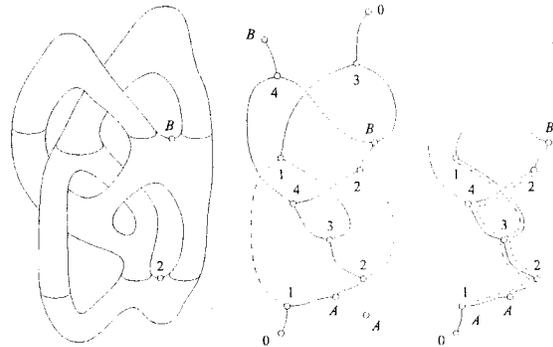

Figure 2: Left: a 2-manifold whose points are mapped to the distance above a horizontal plane. Middle: the Reeb graph in which the critical points of the function appear as degree-1 and degree-3 nodes. The labels indicate the pairing. Right: the tree representing the Reeb graph from slightly above B downwards.

in [4], the Reeb graph has a basis of g loops such that any loop is the sum (modulo 2) of loops in the basis. Each loop has a unique lowest and a unique highest point, referred to as *lo-point* and *hi-point*. We say the lo- and hi-point *span* this loop but note that the same two points may span more than one loop. Each lo-point is an up-fork in the Reeb graph, and each hi-point is a down-fork. We pair each lo-point x with the lowest hi-point y that spans a loop with x. Note that x is also the highest lo-point that spans a loop with y. Indeed, if this were not the case, then we could add the loop spanned by y and x with the loop spanned by y and the lo-point higher than x. The new loop consists of all edges that are either in the first or the second loop but not in both. Its lo-point is x

but it does not contain y so its hi-point must be lower than y, a contradiction. This implies that each lo-point and each hi-point belongs to exactly one pair, giving a total of g pairs between up- and down-forking saddles, as desired.

Algorithm. We construct the Reeb graph of a piecewise linear function on a triangulation with n edges in time $O(n \log n)$ using the algorithm in [4]. It simulates the sweep of a 2-manifold, maintaining the level set as a collection of cyclic lists. The pairing is computed within the same time bound by maintaining a tree during a sequence of lowest common ancestor queries. Specifically, we represent the Reeb graph of $\mathbb{M}_a$ as a forest obtained by gluing branches to each other in a way that eliminates all loops; see Figure 2. The leaves are cyclic lists representing components of $f^{-1}(a)$. We take the following steps at reaching a critical point x, merging two arcs across a degree-2 node whenever one is created:

Case 1: x is a minimum. We add a new tree, consisting of an interior node (the minimum) connected to a single leaf, to the forest.

Case 2: x is an up-forking saddle. We turn the corresponding leaf into an internal node, adding two new leaves as its children.

Case 3: x is a down-forking saddle, connecting leaves u and v. We glue the two downward paths emanating from u and v, ending the gluing at y. If the two paths are disjoint and end at different roots then y is the higher root, a minimum, which we now pair with x (this is a traditional persistence pair). If the downward paths meet then y is the lowest common ancestor (highest in the way we draw trees), an up-fork, which we now pair with x (this is an extended persistence pair).

Case 4: x is a maximum. We pair it with its parent y and remove the joining edge together with the two nodes; y can be either an up-forking saddle, producing a traditional persistence pair, or a minimum, producing an extended persistence pair.

In order to perform these operations efficiently, we use the linking and cutting tree data structure described by Sleator and Tarjan [19]. It decomposes each tree in the forest into vertex-disjoint paths, representing each path using a binary search tree. By maintaining a linking and cutting tree data structure, Cases 1, 2, and 4 can be handled in $O(n \log n)$ overall time. Focusing on Case 3, we spend $O(\log n)$ time to find y using the operations supported by the data structure. The only extra operation we need is gluing two paths of length $k \leq m$ each. As recently shown by Georgiadis, Tarjan and Werneck [12], this gluing operation can be implemented in $O(\log n)$ amortized time. In total, we thus have an algorithm that computes all pairs in time $O(n \log n)$.

Symmetry. The negative function, $-f : \mathbb{M} \to \mathbb{R}$, has the same critical points as f. We claim that it also generates the same pairing.

SYMMETRY LEMMA. Two critical points x and y are paired for f iff they are paired for $-f$.

PROOF. The claim is true for the first minimum, x, and the last maximum, y. Every other pair of f contains at least one saddle. We assume without loss of generality that x is a saddle and that $f(x) < f(y)$. Consider again the sweep of the 2-manifold in the direction of increasing values of f. When we pass $a = f(x)$ we split a cycle in the level set into two. The two cycles belong to the boundary of $\overline{\mathbb{M}}_a$, the set of points with function value a or higher. If the two cycles belong to the same component of $\overline{\mathbb{M}}_a$, such as for the point labeled 2 in Figure 2, then x is a lo-point and y is the lowest hi-point that spans a cycle with x. The claim follows because x is also the highest lo-point that spans a cycle with y. If, on the other hand, the two cycles belong to two different components of $\overline{\mathbb{M}}_a$, such as for the point labeled B in Figure 2, then y is the lower of the two maxima that complete the two components. In the backward sweep (the forward sweep for $-f$), y starts a component that merges into the other, older component at x. Again x and y are also paired for $-f$, which implies the claimed symmetry. □

3 Height and Elevation

In this section, we define the elevation as a real-valued function on a 2-manifold in $\mathbb{R}^3$.

Measuring height and elevation on Earth. Even on Earth, defining the elevation of a point x on the surface is a non-trivial task. Traditionally, it is defined relative to the mean sea level (MSL) in the direction of the measured point. In other words, the *MSL elevation* of a point x is the difference between the distance of x from the center of mass and the distance of the MSL from the center of mass in the direction of x. The difficulty of measuring height in the middle of a continent was overcome by introducing the *geoid*, which is a level surface of the Earth's gravitational potential and roughly approximates the MSL while extending it across land. The *orthometric height* above (or below) the geoid is thus more general but otherwise about the same as the MSL elevation. It is perhaps surprising that the geoid differs significantly from its best ellipsoidal approximation due to non-uniform density of the Earth's crust [11]. Standard global positioning systems (GPS) indeed return the *ellipsoidal height*, which is elevation relative to a standard ellipsoidal representation of the Earth's surface. They also include knowledge of the *geoid height* relative to the ellipsoid and compute the orthometric height of x as its ellipsoidal height minus the geoid height in the direction of x.

A simplifying factor in the discussion of height and elevation on Earth is the existence of a canonical core point, the center of mass. For general surfaces, distance measurements from a fixed center make much less sense. We are interested in this general case, which includes surfaces with non-zero genus for which there is no simple notion of core. Like on Earth, we define the elevation of a point x as the difference between two distances, except we no longer use a reference surface, such as the mean sea level or the geoid, but instead measure relative to a canonically associated other point on the surface. To explain how this works, we give different meanings to the 'height' of a point, which we define for every direction, and its 'elevation', which is the difference between two heights. While height depends on an arbitrarily chosen origin, we will see that elevation is independent of that choice. Indeed, the technical concept of elevation, as introduced shortly, will be similar in spirit to the idea of orthometric height, with the exception that it substitutes the canonical associated point for a globally defined reference surface.

Height, persistence and elevation. Let $\mathbb{M}$ be a smoothly embedded 2-manifold in $\mathbb{R}^3$. We assume that $\mathbb{M}$ is generic but it is too early to say what exactly that should mean. We define the height in a given direction as the signed distance from the plane normal to that direction and passing through the origin. Formally, for every unit vector $\mathbf{u} \in \mathbb{S}^2$, we call $f_{\mathbf{u}}(x) = \langle x, \mathbf{u} \rangle$ the *height* of x in the direction $\mathbf{u}$. This defines a 2-parameter family of height functions,

$$\text{Height} : \mathbb{M} \times \mathbb{S}^2 \to \mathbb{R},$$

where $\text{Height}(x, \mathbf{u}) = f_{\mathbf{u}}(x)$. The height is a Morse function on $\mathbb{M}$ for almost all directions. We pair the critical points of $f_{\mathbf{u}}$ as described in Section 2. Following [7], we define the *persistence* of a critical point as the absolute difference in height to the paired point: $\text{pers}(x) = \text{pers}(y) = |f_{\mathbf{u}}(y) - f_{\mathbf{u}}(x)|$.

Each point $x \in \mathbb{M}$ is critical for exactly two height functions, namely for the ones in the direction of its outer and inner normals: $\mathbf{u} = \pm \mathbf{n}_x$. We proved in Section 2 that the pairs we get for the two opposite directions are the same. Hence, each point $x \in \mathbb{M}$ has a unique persistence, which we use to introduce the *elevation function*,

$$\text{Elevation} : \mathbb{M} \to \mathbb{R},$$

defined by $\text{Elevation}(x) = \text{pers}(x)$. We note that the elevation is invariant under translation and rotation of $\mathbb{M}$ in $\mathbb{R}^3$.

Two-dimensional example. We illustrate the definitions of the height and elevation functions for a smoothly embedded 1-manifold $\mathbb{M}$ in $\mathbb{R}^2$. The critical points of $f_{\mathbf{u}} : \mathbb{M} \to \mathbb{R}$ are the points $x \in \mathbb{M}$ with normal vectors $\mathbf{n}_x = \pm \mathbf{u}$. Figure 3 illustrates a sweep in the vertical upward direction $\mathbf{u}$. Each critical point of $f_{\mathbf{u}}$ starts a component, ends a component by merging it into an older component, or closes the curve. The

critical points that start components get paired with the other critical points.
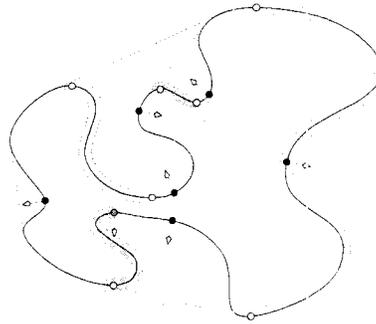

Figure 3: A 1-manifold with marked (white and shaded) critical points of the vertical height function. The shaded strips along the curve connect paired critical points. The black and shaded dots mark two- and one-legged elevation maxima.

The elevation is zero at inflexion points and increases as we move away in either direction. The function may experience a discontinuity at points that share tangent lines with others, such as endpoints of segments that belong to the boundary of the convex hull. The elevation may reach a local maximum at points that either maximize the distance to a shared tangent line or the distance to another critical point in the normal direction. Examples of the first case are the black dots in Figure 3, where the elevation peaks in a non-differentiable manner. An example of the second case is the grey point, where the elevation forms a smooth maximum.

Singular tangencies. The elevation is continuous on $\mathbb{M}$, except possibly at points with singular tangencies. These points correspond to transitional violations of the two genericity conditions of Morse functions. Such violations are unavoidable as Height is a 2-parameter family within which we can transition between Morse functions:

- two critical points may converge and meet at a *birth-death point* where they cancel each other;
- two critical points may *interchange* their positions in the ordering by height, passing a direction at which they share the same height.

The first transition corresponds to an inflexion point of a geodesic on $\mathbb{M}$. Such points are referred to as *flat* or *parabolic*, indicating that their Gaussian curvature is zero. The second transition corresponds to two points $x \neq y$ that share the same tangent plane, $T_x = T_y$. Both types of singularities are forced by varying one degree of freedom and are turned into curves by varying the second degree of freedom. These curves pass through co-dimension two singularities formed by two simultaneous violations of Conditions I and II. There can be two concurrent birth-death points, a birth-death point concurrent with an interchange, or two concurrent interchanges. In each case, the singularity is defined by

4

two pairs of critical points and we get two types each because these pairs may be disjoint or share one of the points. See Table 1 for the features on $\mathbb{M}$ that correspond to the six types of co-dimension two singularities. We can now be more precise about what we mean by a generic 2-manifold, namely that it has a generic family of height functions.

GENERICITY ASSUMPTION A. The 2-parameter family of height functions on $\mathbb{M}$ has no violations of Conditions I and II for Morse functions other than the ones mentioned above (and enumerated in Table 1 below).

Some of these violations will be discussed in more detail later as they can be locations of maximum elevation. Instead of perturbing the 2-manifold, we can enforce the assumptions by perturbing the family of functions, which is easier and suffices for our purposes. A second genericity assumption, referring specifically to the elevation function, will be stated in Section 5.

4 Pedal Surface

In this section, we take a dual view of the height and elevation functions based on a transformation of $\mathbb{M}$ to another surface in $\mathbb{R}^3$. We take this view to help our understanding of the singularities of Height, but it is of course also possible to study them directly using standard results in the field [1, 15].

Pedal function. Recall that T_x is the plane tangent to $\mathbb{M}$ that passes through $x \in \mathbb{M}$. The *pedal p* of x is the orthogonal projection of the origin onto T_x. We write $p = \mathrm{Pedal}(x)$ and obtain a function

$$\mathrm{Pedal} : \mathbb{M} \to \mathbb{R}^3,$$

whose image $\mathbb{P} = \mathrm{Pedal}(\mathbb{M})$ is the *pedal surface* of $\mathbb{M}$ [2]. If the line $0x$ is normal to T_x then $p = x$. More generally, we can construct p by drawing the diameter sphere with center $x/2$ passing through 0 and x. This sphere intersects T_x in a circle with center $(x + p)/2$ that passes through x and $p = \mathrm{Pedal}(x)$. In fact, $\mathbb{P}$ is the evolute of the diameter spheres defined by the origin and the points $x \in \mathbb{M}$, as illustrated in Figure 4. The following three properties are useful in understanding the correspondence between $\mathbb{M}$ and its pedal surface:

- points on $\mathbb{M}$ have parallel and anti-parallel normal vectors iff their images under the pedal function lie on a common line passing through the origin;
- the height of a point $x \in \mathbb{M}$ in the direction of its normal vector is equal to plus or minus the distance of $\mathrm{Pedal}(x)$ from the origin;
- from $p \in \mathbb{P}$ and the angle φ between the vector p and the normal $\mathbf{n}_p$ of $\mathbb{P}$ at p, we can compute the radius ϱ of the corresponding diameter sphere and the preimage x at distance $2\varrho \sin \varphi$ from p in the direction normal to p and $p \times \mathbf{n}_p$.
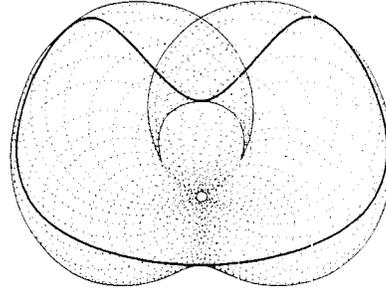

Figure 4: A smoothly embedded (boldface solid) closed curve and the (solid) image of the pedal function constructed as the evolute of the (dotted) diameter circles between the curve and the origin.

The third property implies that the pedal surface determines the 2-manifold.

Co-dimension one singularities. We are interested in singularities of the pedal function as they correspond to directions along which the height function is not generic. For example, a birth-death point of Height corresponds to a cusp point of $\mathbb{P}$. To see this recall that the birth-death points correspond to parabolic points $x \in \mathbb{M}$. A generic geodesic through this point has an inflexion at x, causing the tangent plane to reverse the direction of its rotating motion as we pass through x. Similarly, it causes a sudden reversal of the motion of the image of the origin thus forming a cusp at $\mathrm{Pedal}(x)$. In

Dictionary of Singularities		
$\mathbb{M}$	Height	$\mathbb{P}$
parabolic point double tangency	birth-death (bd) point interchange	cusp xing
Jacobi point triple tangency	2 bd-pts + interchange 3 interchanges bd-pt. + 2 interchanges	dovetail point triple point cusp xing
	2 bd-pts 2 interchanges bd-pt. + interchange	cusp-cusp overpass xing-xing overpass cusp-xing overpass

Table 1: Correspondence between singularities of tangents of the manifold, the 2-parameter family of height functions, and the pedal surface. There are two singularities of co-dimension one: curves of cusps and curves of self-intersections (xings). There are six singularities of co-dimension two, but only two have names in the original domain of tangent planes to $\mathbb{M}$.

contrast, an interchange of Height, which corresponds to a plane tangent to $\mathbb{M}$ in two points, maps to a point of self-intersection (a xing) of $\mathbb{P}$. These two cases exhaust the co-dimension one singularities of Height, which are listed in the upper block of Table 1.

Co-dimension two singularities. There are six types of co-dimension two singularities listed in the lower block of Table 1. Perhaps the most interesting is formed by two concurrent birth-death points that share a critical point, thus forcing a

second interchange. As illustrated in Figure 5, left, the corresponding *dovetail point* in the pedal surface is endpoint of two cusps but also of a self-intersection curve. The sec-

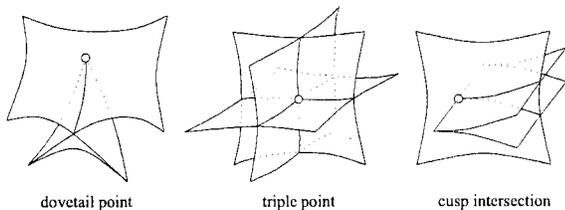dovetail point   triple point   cusp intersection

Figure 5: Left: a portion of the pedal surface in which a self-intersection and two cusps end at a dovetail point. Middle: three sheets of the pedal surface intersecting in a triple point. Right: a cusp intersecting another sheet of the pedal surface.

ond most interesting type is formed by two concurrent interchanges that share a critical point and therefore force a third concurrent interchange of the other two critical points. It corresponds to three self-intersection curves formed by three sheets of $\mathbb{P}$ that intersect in a *triple point*, as shown in Figure 5, middle. Third, we may have a concurrent birth-death point and interchange that share a critical point. As illustrated in Figure 5, right, this corresponds to a cusp curve that passes through another sheet of the pedal surface. There are three parallel types in which the concurrency happens in the same direction **u** but not in space. They correspond to two curves on the pedal surface that cross each other as seen from the origin but do not meet in $\mathbb{R}^3$. As before, a birth-death point corresponds to a cusp curve and an interchange to a curve of self-intersections.

5   Continuity

We are interested in the local maxima of the elevation function, which are the counterparts of mountain peaks and deepest points in the sea. But they are not well defined because the elevation can be discontinuous. We remedy this shortcoming through surgery.

Discontinuities at interchanges. As mentioned in Section 2, the pairs vary continuously as long as the height function varies without passing through birth-death points and interchanges (Conditions I and II). It follows that the elevation is continuous in regions where this is guaranteed. Around a birth-death point, the elevation is necessarily small and goes to zero as we approach the birth-death point. The only remaining possibility for discontinuous elevation is thus at interchanges, which happen when two points share the same tangent plane. As mentioned in Table 1, this corresponds to a point at which the pedal surface intersects itself. Figure 6 shows that discontinuities in the elevation can indeed arise at co-tangent points. We see four points with common vertical normal direction, of which y and z are co-tangent. Consider

a small neighborhood of the vertical direction, **u**, and observe that the critical points vary in neighborhoods of their locations for $f_{\mathbf{u}}$. The critical point near x changes its partner from the right side of y to the left side of z as it varies from left to right in the neighborhood of x. Similarly, the critical point near w changes its partner from the right side of z to the left side of y as it varies from left to right in the neighborhood of w. Since the height difference is the same at the time of the interchange, the elevation at x and w is still continuous. However, it is not continuous at y and at z, which both change their partners, either from x to w or the other way round. Not all interchanges cause discontinuities,
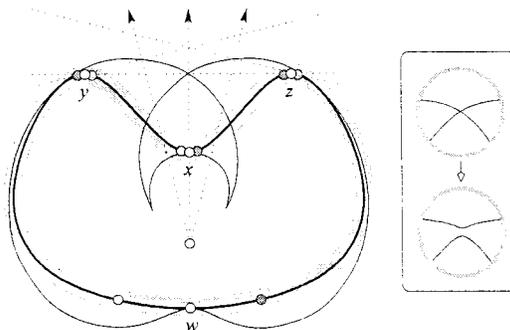

Figure 6: The four white points share the same normal direction, as do the four light shaded and the four dark shaded points. The strips indicate the pairing, which switches when the height function passes through the vertical direction. The insert on the right illustrates the effect of surgery at y and z on the pedal curve.

only those that affect the pairing. These are the interchanges that affect a common topological feature arising during the sweep of $\mathbb{M}$ in the height direction.

Continuity through surgery. We apply surgery to $\mathbb{M}$ to obtain another 2-manifold, $\mathbb{N}$, on which the elevation function is continuous. Specifically, we cut $\mathbb{M}$ along curves at which Elevation : $\mathbb{M} \to \mathbb{R}$ is discontinuous, resulting in a 2-manifold with boundary, $\mathbb{B}$. Then we glue $\mathbb{B}$ along its boundary, making sure that glued points have the same elevation. Formally, we cut by applying the inverse of a surjection $\mathbb{B} \to \mathbb{M}$, and we glue by applying a surjective map from $\mathbb{B}$ to $\mathbb{N}$:

$$\mathbb{M} \xrightarrow{\text{Cut}} \mathbb{B} \xrightarrow{\text{Glue}} \mathbb{N}.$$

As argued above, each boundary curve of $\mathbb{B}$ is defined by an interchange and corresponds to a self-intersection curve (a xing) of the pedal surface. The latter view is perhaps the most direct one in which surgery means cutting along xings and gluing the resulting four sheets in a pairing that resolves the self-intersection. This is illustrated in Figure 6 for a 1-manifold where on the right we see a self-intersection being resolved by cutting the two curves and gluing the upper ends to each other and similarly the lower ends. In the original boldface curve on the left, this operation corresponds to cutting at y and z and gluing the four ends to form two closed

curves: one from y to x to $z = y$ and the other from y to w to $z = y$. As mentioned earlier, not all xings correspond to discontinuities and we perform surgery only on the subset that do. In general, a discontinuity follows a xing until it runs into a dovetail or a triple point. In the former case, the xing and the discontinuity both end. In the latter case, the xing continues through the triple point and the discontinuity may follow, turn, or even branch to other xings passing through the same triple point. There are two possible configurations created by surgery in the neighborhood of a triple point p. Their particular significance in the recognition of local maxima will be discussed shortly. Whatever the situation, the subset of xings along which the elevation is discontinuous together with the gluing pattern across these xings provides a complete picture of how to use surgery to change $\mathbb{P}$ into a new surface, $\mathbb{Q}$. The 2-manifold $\mathbb{N}$ is the one for which this is the pedal surface: $\mathbb{Q} = \text{Pedal}(\mathbb{N})$. That $\mathbb{N}$ is indeed a manifold can be shown by (tedious) examination of all cut-and-glue patterns that may occur. After surgery, we have a continuous function Elevation : $\mathbb{N} \to \mathbb{R}$. Furthermore, we have continuously varying pairs of critical points. To formalize this idea, we introduce a new map

$$\text{Antipode} : \mathbb{N} \to \mathbb{N}$$

that maps a point x to its paired point $y = \text{Antipode}(x)$. The function Antipode is a homeomorphism and its own inverse. We note in passing that we could construct yet another 2-manifold by identifying antipodal points. Each local maximum of the elevation function on this new manifold corresponds to a pair of equally high maxima in $\mathbb{N}$. This construction is the reason we will blur the difference between maxima and antipodal pairs of maxima in the next few sections.

Smoothness. The elevation function on $\mathbb{N}$ is smooth almost everywhere. To describe the violations of smoothness, let $\text{Bd} \, \mathbb{B}$ denote the boundary of the intermediate manifold. Let $B = \text{Glue}(\text{Bd} \, \mathbb{B})$ and define $S = B \cup \text{Antipode}(B)$, the set of points at which the elevation function is not smooth. By Genericity Assumption A, S is a graph, consisting of *nodes* and *arcs*. We have degree-1 and degree-3 nodes that correspond to dovetail points and triple points in the pedal surface, respectively, as well as degree-4 nodes that correspond to overpasses between xings. Each degree-4 node is the crossing of an arc in B and an arc in the antipodal image of B. We think of this construction as a *stratification* of $\mathbb{N}$. Its *strata* are

- the three kinds of nodes of S;
- the open and closed arcs of S;
- the open connected regions in $\mathbb{N} - S$.

Figure 7 illustrates the construction by showing how such a stratification may look like. Because of Genericity Assumption A, the restriction of the elevation function to a stratum is smooth and a Morse function if we remove the
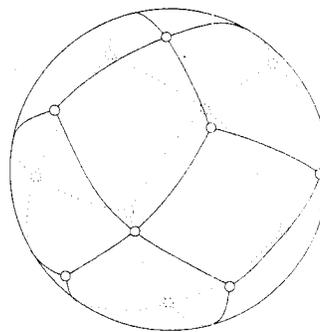

Figure 7: Stratification of the 2-sphere obtained by overlaying a spherical tetrahedron with its antipodal image. The (shaded) degree-4 nodes are crossings between B and its antipodal image.

points where it is zero. Let $\mathbb{N}_0$ be $\mathbb{N}$ minus the points x with $\text{Elevation}(x) = 0$. The set S defines a stratification of $\mathbb{N}_0$ whose strata are the components of the earlier strata after removing the zero elevation points. A *stratified Morse function* on $\mathbb{N}_0$ satisfies

(i) its restriction to a stratum is a Morse function;

(ii) the limit of the derivative as we approach a critical point on the boundary of a stratum is non-zero;

see [13]. We are now ready to state our second genericity assumption.

GENERICITY ASSUMPTION B. The elevation function restricted to $\mathbb{N}_0$ is a stratified Morse function with strata defined by S.

The implication of this assumption will become clear after we enumerate the generic types of local maxima of the elevation function in next section. For example, it implies that the sphere is not generic.

6 Elevation Maxima

In this section, we enumerate the generic types of local maxima of the elevation function. They come in pairs in $\mathbb{N}$ which, by inverse surgery, form multi-legged creatures in $\mathbb{M}$.

Classification of local maxima. Depending on its location, a point $x \in \mathbb{N}$ can have one, two, or three preimages under surgery. We call this number its *multiplicity*, $\mu(x)$. Specifically, x has multiplicity three if it is a node of the graph B, it has multiplicity two if it lies on an arc of B, and it has multiplicity one otherwise. Degree-4 nodes in the stratification of $\mathbb{N}$ correspond to antipodal pairs of points with multiplicity two each. Let now $x \in \mathbb{N}$ be a local maximum of the elevation function. We know that x is not a parabolic point of $\mathbb{M}$, else its elevation would be zero. This simple observation eliminates five of the eight singularities in Table 1. Furthermore, Genericity Assumption A on 2-manifold $\mathbb{M}$ implies

that a multiplicity three point can only be paired with a multiplicity one point. This leaves the following four possible types of local maxima x:

$$\left.\begin{array}{l} \textit{one-legged} \\ \textit{two-legged} \\ \textit{three-legged} \\ \textit{four-legged} \end{array}\right\} \text{ if } \left\{\begin{array}{l} \mu(x) = \mu(y) = 1, \\ \mu(x) = 1 \text{ and } \mu(y) = 2, \\ \mu(x) = 1 \text{ and } \mu(y) = 3, \\ \mu(x) = \mu(y) = 2, \end{array}\right.$$

where $y = \text{Antipode}(x)$; see Figure 1. We sometimes call the preimages of x the *heads* and those of y the *feet* of the maximum. The most exotic of the four types is perhaps the four-legged maximum, which corresponds to an overpass of two xings in the pedal surface or, equivalently, a degree-4 node in the stratification. The image of x under Pedal lies on one xing and the image of y lies on the other. Both points have two preimages under surgery, which makes for a complete bipartite graph with two heads, two feet, and four legs.

Neighborhood patterns. It is instructive to look at the local neighborhood of a maximum x in $\mathbb{M}$. Most interesting is the three-legged type, with feet y_1, y_2, y_3. A small perturbation of the normal direction can change the ambiguous pairing of x with all three to an unambiguous pairing of a point in the neighborhood of x with a point in the neighborhood of one of the feet. To study these pairings, we map a neighborhood of x to a neighborhood of $\mathbf{n}_x$ using the Gauss map. Assuming x is not a parabolic point of $\mathbb{M}$, the Gauss map is locally homeomorphic. The unambiguous pairings are indicated by labeling every direction in the neighborhood of $\mathbf{n}_x$ with the index of the corresponding foot, as in Figure 8. The three curves passing through $\mathbf{n}_x$ correspond to the
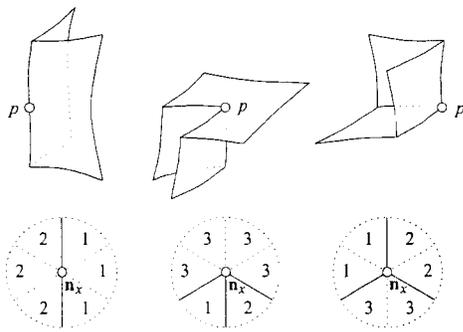

Figure 8: Top row: the three sheets of $\mathbb{Q}$ after possible ways of cutting and gluing the neighborhood of a triple point p in $\mathbb{P}$. Bottom row: the corresponding pairing patterns in the neighborhood of $\mathbf{n}_x$, each belonging to one of the three points paired with the three feet. The (shaded) Mercedes star is necessary for a three-legged maximum.

three xings passing through the triple point $p \in \mathbb{P}$, which is the simultaneous image of the three feet. They decompose the neighborhood into six slices corresponding to the six permutations of the feet. The labeling indicates the pairing and reflects the surgery at these feet and, equivalently, at

the corresponding triple point in the pedal surface. Only the rightmost pattern in Figure 8 corresponds to a maximum, the reason of which will become clear shortly.

The neighborhood pictures for the remaining three types of maxima are simpler. For a one-legged maximum we have an undivided disk, which requires no surgery. For a two- or four-legged maximum we have a disk divided into two halves and there is only one way to do the surgery.

Necessary projection conditions. Let $x \in \mathbb{N}$ be a maximum of Elevation : $\mathbb{N} \to \mathbb{R}$ with $y = \text{Antipode}(x)$. Recall that the corresponding heads $x_i \in \mathbb{M}$ all have the same tangent plane. Similarly, the corresponding feet $y_j \in \mathbb{M}$ have the same tangent plane, which is parallel to the one of the x_i. We formulate additional conditions necessarily satisfied by maxima of the elevation function.

PROJECTION CONDITIONS. The point x is a maximum of the elevation function only if

#legs = 1: $\mathbf{n}_x$ is parallel or anti-parallel to $y - x$;

#legs = 2: $\mathbf{n}_x$, $y_1 - x$, $y_2 - x$ are linearly dependent and the orthogonal projection of x onto the line of the two feet lies between y_1 and y_2;

#legs = 3: the orthogonal projection of x onto the plane of the three feet lies inside the triangle spanned by y_1, y_2, y_3;

#legs = 4: the orthogonal projections of the line segments $x_1 x_2$ and $y_1 y_2$ onto a plane parallel to both have a non-empty intersection.

In summary, x is a local maximum only if $\mathbf{n}_x$ or $-\mathbf{n}_x$ is a positive linear combination of the vectors $y_j - x_i$.

It is easy to verify the analogous conditions for a smooth function on a 1-manifold. There is only one type of pair, consisting of a minimum x and a maximum y. In the one-legged case, y has to be right above x, in the direction $\mathbf{n}_x$, else we have the configuration depicted in Figure 9, left, and we can increase the height difference by moving the normal direction toward the direction of $y - x$. To see this, consider the osculating circles touching the 1-manifold in points x and y. Simultaneously roll a line about each circle, keeping the two lines parallel at all times. The distance between the lines, which is the height difference between the points at which they touch the circles, is a local maximum iff the difference vector between these two points is normal to the rolling lines. If this is not the case, there is a small motion of the normal direction that increases the height difference, as indicated in Figure 9, left. We remark that the argument does not depend on x being a minimum and y being a maximum and also works for other combinations of critical point types. In the two-legged case, the orthogonal projection of x onto the line of y_1 and y_2 has to lie between the two points. Else we have the configuration depicted in Figure 9, right, and we can move the normal direction toward the directions of $y_1 - x$ and $y_2 - x$, which lie on the same side of $\mathbf{n}_x$, thus increasing

the height difference for both pairs. One of the (perturbed) pairs remains paired by persistence, contradicting that x is an elevation maximum.
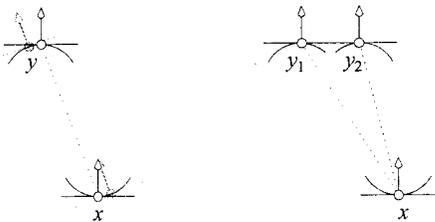

Figure 9: A one-legged configuration on the left and a two-legged configuration on the right. In both configurations, we can increase the height difference by tilting the normal direction towards the left.

The argument for a 2-manifold is essentially the same, except that we have more cases and points with 2-dimensional neighborhoods. An interesting case is the three-legged maximum, formed by x and three feet, y_1, y_2, y_3. Call a plane *anchored* if it passes through x and contains the direction $\mathbf{n}_x$. To establish that the orthogonal projection of x onto the plane of the three feet lies inside the triangle they span, we assume the opposite and take a plane anchored at x such that all three feet lie on one side of the plane. Take a second plane anchored at x that intersects the first plane at a right angle and project the configuration into that plane, drawing the sectional curve parallel to the second plane for each critical point. The picture is like Figure 9, right, except we now have three points y_j, and we use the 1-dimensional argument to contradict the maximality of this configuration.

As mentioned earlier, we have an additional necessary condition for the three-legged case, namely that the neighborhood pattern at $\mathbf{n}_x$ be a Mercedes star. Indeed, suppose it is not, then the pattern contains a straight line through the center along which we read the same label j at all points. In other words, there is a plane anchored at x such that for all normal directions close to $\mathbf{n}_x$ that lie in this plane the persistence algorithm pairs (a point in the neighborhood of) x with (a point in the neighborhood of) y_j. If there is only one such plane, like in the pattern on the left in Figure 8, genericity contradicts maximality. Otherwise, we get two such anchored planes and use the 1-dimensional argument within these planes to show that $y_j - x$ is parallel or anti-parallel to $\mathbf{n}_x$. But x cannot simultaneously be part of a one-legged and a three-legged maximum without contradicting the condition for the derivative of f stated in the Genericity Assumption B.

7 Algorithm

In this section, we describe an algorithm for constructing all points with locally maximum elevation. The input is a piecewise linear 2-manifold embedded in $\mathbb{R}^3$. The running time of the algorithm is polynomial in the number of edges.

Smooth vs. piecewise linear. We consider the case in which the input is a two-dimensional simplicial complex K in $\mathbb{R}^3$. This data violates some of the assumptions we used in our mathematical considerations. This causes difficulties which, with some effort, can be overcome. For example, it makes sense to require that K be a 2-manifold but not that it be smoothly embedded. The 2-parameter family of height functions is well-defined and continuous but not smooth. The definition of the elevation function is more delicate as it makes reference to point pairs in all possible directions. For any given direction, we get a well-defined collection of pairs, but how can we be sure that the pairs for different directions are consistent? A related difficulty is rooted in the fact that a vertex in K is usually critical for many directions and it may be paired with different vertices in different directions. To rationalize this phenomenon, we think of K as the limit of an infinite series of smoothly embedded 2-manifolds. A vertex of K gets resolved into a small patch with a two-dimensional variety of normal directions. As the patch shrinks toward the vertex, the variety of normal directions approaches a limit region on the 2-sphere. For different directions in this limit region, the corresponding points on the patch may be paired with points from different other patches. It thus seems natural that, in the limit, a vertex would be paired with more than one other point.

Next, we introduce a combinatorial notion of the variety of normal directions. Let σ be a simplex in K and recall that the link of σ is empty if σ is a triangle, it is a pair of vertices if σ is an edge, and it is a cycle of vertices and edges if σ is a vertex. Let $\mathbf{u} \in \mathbb{S}^2$ be a direction such that $\langle \mathbf{u}, x \rangle = \langle \mathbf{u}, z \rangle$ for all points $x, z \in \sigma$. Given $\mathbf{u}$, the *upper link* of σ consists of all points y in the link for which $\langle \mathbf{u}, y \rangle > \langle \mathbf{u}, x \rangle$. We say x is *critical* for the height function in the direction $\mathbf{u}$ if the complement of the upper link is not contractible. For example, the empty lower link of a minimum and the complete circle of a maximum are both not contractible. Let $\mathbf{N}(x) \subseteq \mathbb{S}^2$ be the set of directions along which x is critical. Generically, the set $\mathbf{N}$ for a point inside a triangle of K is a pair of opposite points on the sphere, that for a point on an edge is a pair of opposite open great-circle arcs, and that for a vertex is a pair of opposite open spherical polygons. Here, the word 'generic' applies to a simplicial complex in $\mathbb{R}^3$, where it simply means that the vertices are in general position. Computationally, this assumption can be simulated by a symbolic perturbation [9]. We write $\mathbf{N}(x, y, \ldots)$ for the common intersection of the sets $\mathbf{N}$ of x, y, and so on.

Finite candidate sets. Given a candidate for a maximum, we can use the extended persistence algorithm to decide whether or not it really is a maximum. More specifically, we need a point x and a direction $\mathbf{u}$ along which the sweep defining the pairing proceeds. The details of this decision algorithm will be discussed shortly. We use the Projection Conditions, which are necessary for local maxima, to get four kinds of candidates:

#legs = 1: pairs of points x and y on K with the direction $(y - x)/\|y - x\|$ contained in $\mathbf{N}(x, y)$;

#legs = 2: triplets of points x, y_1, y_2 such that the orthogonal projection z of x onto the line of y_1 and y_2 lies between the two points and the direction $(z - x)/\|z - x\|$ is contained in $\mathbf{N}(x, y_1, y_2)$;

#legs = 3: quadruplets of points x, y_1, y_2, y_3 such that the orthogonal projection z of x onto the plane of y_1, y_2, y_3 lies inside the triangle and the direction $(z - x)/\|z - x\|$ is contained in $\mathbf{N}(x, y_1, y_2, y_3)$;

#legs = 4: quadruplets of points x_1, x_2, y_1, y_2 such that the shortest line segment zw connecting the lines of x_1, x_2 and y_1, y_2 also connects the two line segments and the direction $(z - w)/\|z - w\|$ is contained in $\mathbf{N}(x_1, x_2, y_1, y_2)$.

With the assumption of a generic simplicial complex K, we get a finite set of candidates of each kind. Since this might not be entirely obvious, we discuss the one-legged case in some detail. Let σ and τ be two simplices and x and y points in their interiors. For a generic K, the intersection of normal directions, $\mathbf{N}(x, y)$, is non-empty only if one of the two simplices is a vertex or both are edges. If $x = \sigma$ is a vertex then y is necessarily the orthogonal projection of x onto τ, which may or may not exist. If σ and τ are both edges then xy is necessarily the line segment connecting σ and τ and forming a right angle with both, which again may or may not exist. In the end, we get a set of $O(n^2)$ candidate pairs x and y, where n is the number of edges in K. For the two-legged case, we get $O(n^3)$ candidates, each a triplet of vertices or a pair of vertices together with a point on an edge. For the three- and four-legged cases, we get $O(n^4)$ candidates, each a quadruplet of vertices, giving a total of $O(n^4)$ candidates.

Verifying candidates. Let $x, y \in \mathbb{N}$ be a pair of points whose heads and feet all have parallel or anti-parallel normal directions. In the smooth case, two the necessary conditions for x and y to define an elevation maximum are:

(a) the Projection Conditions of Section 6;

(b) the requirement that $y = \mathrm{Antipode}(x)$.

We subsume the Mercedes star property in (b) since it depends on the antipodality map or, equivalently, on the pairing by extended persistence. If x and y are indeed a maximum pair, then they also satisfy a constraint on their curvature that ensures a local decrease of the height difference in a neighborhood of $\mathbf{n}_x$. In the piecewise linear case, this curvature constraint is redundant because the curvature concentrates at the vertices and edges, and thus we have only (a) and (b). We have seen above how to translate (a) to the piecewise linear case. It remains to test (b), which reduces to answering a constant number of *antipodality queries*: given a direction $\mathbf{u}$ and a critical point x of $f_{\mathbf{u}}$, find the paired critical point y. This is part of what the algorithm described

in Section 2 computes while sweeping K in the direction $\mathbf{u}$. More precisely, the algorithm computes one of the possible pairs, if applied in non-generic directions in which two or more vertices share the same height. Most of our candidates generate non-generic directions, and we cope with this situation by running the algorithm several times, namely once for each combination of permutations of the heads and of the feet. Each combination corresponds to a generic direction that is infinitesimally close to the non-generic direction. The largest number of combinations is six, which we get for three-legged maxima. This is also how we decide the Mercedes star property: each foot is answer to two of the six antipodality queries, giving rise to the pattern in Figure 8. Letting n be the number of edges, the algorithm takes time $O(n \log n)$ to answer an antipodality query. Since we have $O(n^4)$ candidates to test, this amounts to a total running time of $O(n^5 \log n)$.

8 Discussion

The main contribution of this paper is the definition of elevation as a real-valued function on a 2-manifold embedded in $\mathbb{R}^3$ and the computation of its local maxima. We have implemented the algorithm and used it to predict docking configurations for protein structures [20]. Given the surfaces of two such structures, we compute their elevation maxima and align them to suggest docking configurations to be evaluated. Each alignment is a rigid motion, a point in six-dimensional space of translations composed with rotations, and the algorithm may be interpreted as a way to sample this space. Since the elevation function also provides estimates for the direction and size of features, we observed that the generated sampling is significantly more efficient (smaller) than those of other geometric docking algorithms. However, the obtained configurations are not very accurate and need to be refined by a local improvement algorithm. An apparent advantage of our algorithm over others is that it finds good coarse alignments also for proteins whose shapes are known only approximately, for example if they undergo small deformations when they form complexes.

It would be worth exploring extensions of our results to manifolds with boundary and to manifolds of dimension three or higher. A crucial first step will have to be the generalization of the concept of extended persistence to these more general topological spaces. The algorithm presented in Section 7 enumerates all local maxima of the elevation function, without computing the elevation function itself, other than at a collection of candidate points. This approach is suggested by the ambiguities that arise in the definition of the elevation function for piecewise linear data. Unfortunately, it implies the fairly high running time of $O(n^5 \log n)$ in the worst case. Can the maxima be enumerated more efficiently than that? Is there an algorithm that enumerates all maxima above some elevation threshold without computing the maxima below the threshold?

References

[1] V. I. ARNOLD. *Catastrophe Theory.* Springer-Verlag, Berlin, Germany, 1984.

[2] J. W. BRUCE AND P. J. GIBLIN. *Curves and Singularities.* Second edition, Cambridge Univ. Press, England, 1992.

[3] F. CAZALS, F. CHAZAL AND T. LEWINER. Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. *In* "Proc. 19th Ann. Sympos. Comput. Geom., 2003", 351–360.

[4] K. COLE-MCLAUGHLIN, H. EDELSBRUNNER, J. HARER, V. NATARAJAN AND V. PASCUCCI. Loops in Reeb graphs of 2-manifolds. *Discrete Comput. Geom.* **32** (2004), 231–244.

[5] M. L. CONNOLLY. Shape complementarity at the hemoglobin albl subunit interface. *Biopolymers* **25** (1986), 1229–1247.

[6] F. H. C. CRICK. The packing of alpha-helices: simple coiled coils. *Acta Crystallogr.* **6** (1953), 689–697.

[7] H. EDELSBRUNNER, J. HARER AND A. ZOMORODIAN. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete Comput. Geom.* **30** (2003), 87–107.

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

[9] H. EDELSBRUNNER AND E. P. MÜCKE. Simulation of Simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104.

[10] A. H. ELCOCK, D. SEPT AND J. A. MCCAMMON. Computer simulation of protein-protein interactions. *J. Phys. Chem.* **105** (2001), 1504–1518.

[11] W. FRACZEK. Mean sea level, GPS, and the geoid. *ArcUsers Online*, ERSI Web Sites: www.esri.com/news/arcuser/0703/summer-2003.html, 2003.

[12] L. GEORGIADIS, R. E. TARJAN AND R. F. WERNECK. Design of data structures for mergeable trees. Manuscript, Comput. Sci. Dept., Princeton Univ., New Jersey, 2005.

[13] M. GORESKY AND R. MACPHEARSON. *Stratified Morse Theory.* Springer-Verlag, Heidelberg, Germany, 1988.

[14] I. HALPERIN, B. MA, H. WOLFSON AND R. NUSSINOV. Principles of docking: an overview of search algorithms and a guide to scoring functions. *Proteins* **47** (2002), 409–443.

[15] A. HATCHER AND J. WAGONER. *Pseudo-Isotopies of Compact Manifolds.* Société Mathématique de France, 1973.

[16] J. JANIN AND S. J. WODAK. The structural basis of macromolecular recognition. *Adv. Protein Chem.* **61** (2002), 9–73.

[17] J. MILNOR. *Morse Theory.* Princeton Univ. Press, New Jersey, 1963.

[18] J. C. MITCHELL, R. KERR AND L. F. TEN EYCK. Rapid atomic density measures for molecular shape characterization. *J. Mol. Graph. Model.* **19** (2001), 324–329.

[19] D. D. SLEATOR AND R. E. TARJAN. A data structure for dynamic trees. *J. Comput. Sys. Sci.* **26** (1983), 362–391.

[20] Y. WANG, P. K. AGARWAL, P. BROWN, H. EDELSBRUNNER AND J. RUDOLPH. Coarse and reliable geometric alignment for protein docking. *In* "Proc. Pacific Sympos. Biocomputing 2005", World Scientific, Singapore, 64–75.