Approximation by skin surfaces

Nico Kruithof* and Gert Vegter
Department of Mathematics and Computing Science
University of Groningen, The Netherlands
P.O. Box 800
9700 AV Groningen, The Netherlands
{nico,gert}@cs.rug.nl

ABSTRACT

We present a method to approximate a simple, regular C^2 surface W in $\mathbb{R}^3$ by a (tangent continuous) skin surface S. The input of our algorithm is a set of approximate W-maximal balls, where the boundary of the union of these balls is homeomorphic to W. By generating hyperboloid and spherical patches over the intersection curves of the balls the algorithm determines a one-parameter family of skin-surfaces, where a parameter controls the size of the patches. The skin surface S is homeomorphic to W, and the approximate W-maximal balls in the input set are also S-maximal. The Hausdorff distance between the regions enclosed by the input surface W and the approximating skin surface S depends linearly on a parameter related to the sampling density of the approximate W-maximal balls.

Categories and Subject Descriptors

G.1.2 [**Numerical analysis**]: Approximation—*Approximation of surfaces and contours*; I.3.5 [**Computer graphics**]: Computational Geometry and Object Modeling—*Curve, surface, solid, and object representations*; I.3.5 [**Computer graphics**]: Computational Geometry and Object Modeling—*Geometric algorithms, languages, and systems*

General Terms

Algorithms, Design, Theory

Keywords

Smooth surface approximation, skin surfaces, guaranteed topology, Voronoi diagram, Power diagram, Regular triangulation, Medial Axis Transform

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

1. INTRODUCTION

We consider the problem of approximating a simple, regular smooth (C^2) closed surface in $\mathbb{R}^3$ by a skin surface S. Skin surfaces, introduced by Edelsbrunner in [7], are mainly used for modeling large molecules in biological computing in particular, the Van der Waals surface of a molecule. Each atom in the molecule is represented by a sphere and atoms that lie close to each other are connected by patches. These patches make the transition between the atoms tangent continuous. A skin surface is parametrized by a set of weighted points (balls) and a shrink factor. If the shrink factor equals one, the surface is just the boundary of the union of the maximal balls. If the shrink factor decreases, the skin surface becomes tangent continuous, due to the appearance of hyperboloid and spherical patches connecting the balls. An example in $2D$ of a curve, reconstructed using this approach, is drawn in Figure 1c. In Figure 2 some skin surfaces are drawn for different values of this shrink factor. Other nice properties of these skin surfaces are fast visualization, tangent continuity and ease of construction and morphing.

The set of weighted points defining the skin surface S approximating W is a finite sample of the medial axis transform of W, i.e., a finite set of W-maximal balls. Recall that an empty ball of W is a ball enclosed by W. A W-maximal ball is an empty ball of W not contained in any other empty ball of W. The *medial axis transform* of W is the set of W-maximal balls. The *medial axis* M of W is the closure of the set of centers of the maximal balls of W, and can be seen as the skeleton of the surface. In fact, the medial axis is a deformation retract of the region bounded by the surface.

A sufficiently dense finite subset of the medial axis transform, or, rather, the boundary of the union of the corresponding maximal balls, forms a good approximation of the surface. Figure 1b illustrates this observation for a curve in the plane. Obviously, this approximation is not tangent continuous. Our algorithm reconstructs a tangent continuous surface from a sample of (approximate) maximal balls by adding smooth patches over the points of intersection. By controlling the size of these patches, we guarantee that the approximation is homeomorphic to the union of the maximal balls. Furthermore, as we increase the sampling density of the set of (approximate) maximal balls, the maximal distance between the regions enclosed by W and the approximating skin surface tends to zero.

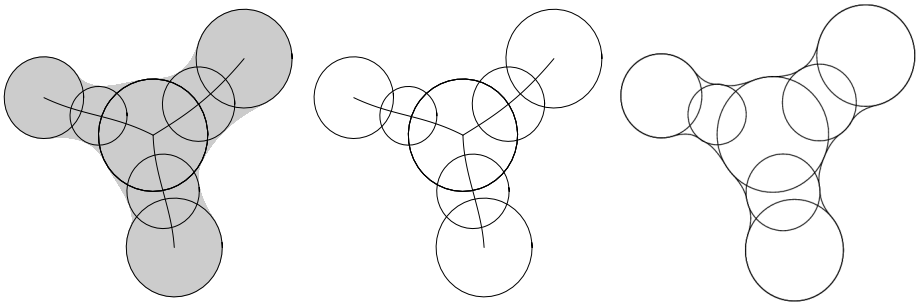

(a) (b) (c)

Figure 1: The approximation of a curve in the plane. (a) The medial axis and some maximal circles. (b) A finite set of maximal circles is the input of our algorithm. (c) The output is a skin curve approximating the input curve in Figure 1a.

In Section 2 we review the theory of skin surfaces as presented in [7], and describe the method to subdivide these surfaces into patches of degree two (patches of hyperboloids and balls). With the language of skin surfaces in place, we give a more precise specification of the algorithm in Section 3, and state the main result on its complexity: the skin surface can be computed in $O(N^2 \log N)$ time and $O(N^2)$ space, where N is the number of approximate maximal balls in the input set. A key ingredient of the algorithm is the construction of the weighted Voronoi Diagram of a set of balls with radii growing with the same multiplicative factor. In Section 4 we present an algorithm that maintains this diagram as the radii are growing. In Section 5 we show how the algorithm satisfies the constraint that the balls in the input set are also S-maximal. Finally, Section 6 derives a bound on the Hausdorff distance between the regions enclosed by the input surface W and the skin surface S in terms of a parameter related to the sampling density of the set of (approximate) maximal balls.

Related work.

Amenta et al. [1] show that the medial axis transform can effectively be approximated by a sufficiently dense sample from the set of maximal balls. If the sample is sufficiently dense, the boundary of the union of the balls is homeomorphic to the original surface. Then they use an approximation of the medial axis transform to compute the power crust, which approximates the original surface by piecewise linear patches.

Another method for visualising molecules uses Connolly surfaces; See, e.g., [3]. For further reading about skin surfaces we refer to [7], and to [5, 6] for background on morphing skin surfaces. In [4] a method is described to triangulate a skin surface with shrink factor 1/2 and to maintain this triangulation under a specific growth model. In this model the radius of each input ball is changed by adding a constant to its square. Our growth model multiplies the squared radius with a constant and the shrink factor is not restricted to 1/2, therefore this triangulation can not be used to triangulate our approximating surfaces.

Boissonnat and Cazals [2] also reconstruct a tangent continuous surface from a sufficiently dense finite sample of points on a surface. Their method uses a different approach using natural neighbor interpolation.

2. SKIN SURFACES

The definitions in this section form a summary of [7]. For further reading on skin surfaces we refer to this article.

A skin surface is defined in terms of a finite set of balls (weighted points) and a shrink factor s, with $0 < s \leq 1$. For $s = 1$, the skin surface is the boundary of the union of balls. For smaller s, the radii of the balls in the input set are shrunken, and hyperboloid and spherical patches are generated between adjacent balls, making the skin surface tangent continuous.

All definitions are given for skin surfaces in $3D$, but can easily be transformed into the definition of skin curves in $2D$. All images show skin curves in the plane.

The space of weighted points.

A $3D$ weighted point $\hat{p}$ is a pair $\hat{p} = (p, \mathrm{P})$ of location $p \in \mathbb{R}^3$ and weight $\mathrm{P} \in \mathbb{R}$. Balls correspond to weighted points, with location at the center of the ball and weight equal to the square of the radius. This definition is slightly different from [7], in the sense that there a weighted point is associated with the sphere bounding the ball. On the set of weighted points we define a 'distance' function:

$$\pi(\hat{p}, \hat{q}) = \| p - q \|^2 - \mathrm{P} - \mathrm{Q}, \qquad (1)$$

for $\hat{p} = (p, \mathrm{P})$ and $\hat{q} = (q, \mathrm{Q})$. Two points with zero distance are by definition *orthogonal*. We will call a point with zero weight an *unweighted point*. The distance from a weighted point $\hat{p}$ to an unweighted point x, follows from Equation (1), by taking $q = x$ and $\mathrm{Q} = 0$. All unweighted points with non-positive distance to $\hat{p}$ form the ball with center p and radius $\sqrt{\mathrm{P}}$. We will make no distinction between the weighted point and this ball. An *orthosphere* of a set of weighted points $\mathcal{P}$ is, by definition, a sphere orthogonal to each of the weighted points in $\mathcal{P}$.

The space of weighted points inherits a vector space structure from $\mathbb{R}^4$ via the bijective map $\Pi : \mathbb{R}^3 \times \mathbb{R} \to \mathbb{R}^4$, defined by

$$\Pi(\hat{p}) = (\xi_1, \xi_2, \xi_3, |p|^2 - \mathrm{P}), \text{ with } p = (\xi_1, \xi_2, \xi_3).$$

The notion of *shrinking* and growing of weighted points can be defined in terms of this vector space structure. Given a shrink factor $s \in \mathbb{R}$, we associate with a weighted point $\hat{p}$ the shrunken weighted point $\hat{p}^s$ defined as follows. Let

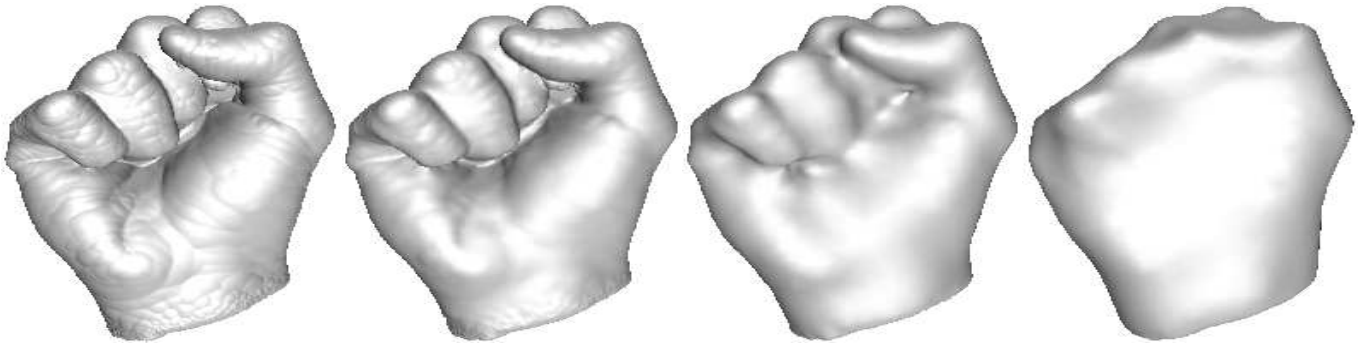

Figure 2: The approximating skin surface of a set of balls forming a hand. The patches between the maximal balls become larger as the shrink factor is decreased. The shrink factors (from left to right) are 0.95, 0.8, 0.5 and 0.2.

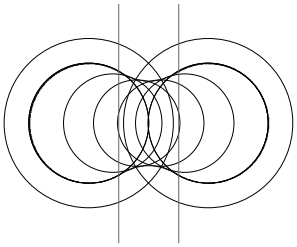

Figure 3: The skin curve of two weighted points (the two larger circles). The smaller circles form a subset of the shrunken convex hull of the input points. The boundary of the shrunken convex hull forms the skin curve.

$\hat{p}' = (p, 0)$. Then

$$\hat{p}^s = s \cdot \hat{p} + (1 - s) \cdot \hat{p}'$$

Using simple calculus it follows that:

$$\begin{aligned} \hat{p}^s &= \Pi^{-1}(s \cdot \Pi(\hat{p}) + (1 - s) \cdot \Pi(\hat{p}')) \\ &= (p, s \cdot \mathrm{p}) \end{aligned}$$

If $\mathcal{P}$ is a set of weighted points, we denote by $\mathcal{P}^s$ the set obtained by shrinking every point of $\mathcal{P}$ by a factor s. The body of a skin surface $\mathtt{bdy}^s\mathcal{P}$ and skin surface $\mathtt{skn}^s\mathcal{P}$, associated with a set $\mathcal{P}$ of weighted points, is defined by

$$\mathtt{bdy}^s\mathcal{P} = \bigcup \mathtt{conv}(\mathcal{P})^s \quad (2)$$

$$\mathtt{skn}^s\mathcal{P} = \partial\, \mathtt{bdy}^s\mathcal{P}. \quad (3)$$

Here $\mathtt{conv}(\mathcal{P}) \subset \mathbb{R}^3 \times \mathbb{R}$ is the convex hull – with respect to the vector space structure inherited under Π – of a set of weighted points $\mathcal{P}$, whereas ∂ denotes the boundary – in three space – of the balls. For a skin curve in $2D$ associated with two weighted points: see Figure 3.

Two sets are homeomorphic (or: have the same topology) if there exists a continuous map with a continuous inverse between the sets, see [8, Ch. 28]. One important feature of a skin surfaces is that varying the shrink factor does not change the topology of skin surfaces. For $s = 1$, the body of the skin surface is the union of the balls in the input set, therefore the body of a skin surface is homeomorphic to the union of the balls in the input set.

Quadratic patches of a skin surface.

One of the main features of skin surfaces is that they can be effectively subdivided into patches of degree two, more precisely, into parts of spheres and hyperboloids. In this section we will show how this is done. For the proof that these patches form the corresponding skin surface, we refer to [7].

The mixed complex, which subdivides the skin surface into patches of degree two, is an intermediate structure between the weighted Delaunay triangulation (or: regular triangulation) and its dual, viz. the weighted Voronoi diagram (or: the power diagram).

The weighted Voronoi diagram.

The *weighted Voronoi cell*, Voronoi cell for short, of a weighted point $\hat{p} \in \mathcal{P}$ consists of all points that are closer to $\hat{p}$ than to any other point in $\mathcal{P}$. Formally:

$$V_{\hat{p}} = \{x \in \mathbb{R}^3 | \pi(\hat{p}, x) \leq \pi(\hat{q}, x), \text{ for all } \hat{q} \in \mathcal{P}\}.$$

A weighted Voronoi cell is a convex polyhedral region, possibly unbounded, or empty. Furthermore, a weighted Voronoi cell $V_{\hat{p}}$ does not necessarily contain p.

For a subset $\mathcal{X} \subseteq \mathcal{P}$ we have:

$$\nu_{\mathcal{X}} = \bigcap_{\hat{p} \in \mathcal{X}} V_{\hat{p}}$$

which we will call a *weighted Voronoi ℓ-cell* ($\ell = \dim \nu_{\mathcal{X}} = 4 - |\mathcal{X}|$), if it is not empty and $\mathcal{X}$ is in general position. Then it also follows that $|\mathcal{X}| \leq 4$. A set is in general position if no 5 weighted points are equidistant to some point and no $k + 2$ points lie on a common k-flat for $k = 0, 1, 2$. In the non-degenerate case, every 0-cell is a point, which is the intersection of four adjacent cells. A 1-cell is a line segment, possibly unbounded. And a facet (2-cell) is the boundary of two 3-cells. The 2-cell is a subset of the set of unweighted points equidistant to the two weighted points associated with the 3-cells. A 3-cell is a nonempty Voronoi-cell of some $\hat{p} \in \mathcal{P}$.

The *weighted Voronoi diagram* is the subdivision of the 3-space generated by the non-empty Voronoi cells:

$$\mathrm{Vor}\,\mathcal{P} = \{\nu_{\mathcal{X}} | \mathcal{X} \subseteq \mathcal{P} \wedge \nu_{\mathcal{X}} \neq \emptyset\}.$$

The weighted Delaunay triangulation.

The weighted Delaunay triangulation, or regular triangulation, is dual to the weighted Voronoi diagram, and is defined in terms of weighted Delaunay cells. A *weighted Delaunay* $3-\ell$*-cell* exists for every weighted Voronoi ℓ-cell $\nu_{\mathcal{X}} \in \text{Vor } \mathcal{P}$ and is the convex hull of the centers of the points in $\mathcal{X}$, $\delta_{\mathcal{X}} = \texttt{conv}(\{p|\hat{p} \in \mathcal{X}\})$. The *weighted Delaunay triangulation* is the set of all these cells:

$$\text{Del}\,\mathcal{P} = \{\delta_{\mathcal{X}} | \nu_{\mathcal{X}} \in \text{Vor } \mathcal{P}\}.$$

Note that $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ for $\nu_{\mathcal{X}} \neq \emptyset$ can be disjoint, but their affine hulls always intersect in one point. We will call this point the *focus* f of $\mathcal{X}$. Since the focus lies on the affine hull of the weighted Voronoi cell, it follows that $\pi(\hat{p}, f) = \pi(\hat{q}, f)$ for all $\hat{p}, \hat{q} \in \mathcal{X}$.

The mixed complex.

The *mixed complex*, associated with a scalar s, with $0 \leq s \leq 1$, is obtained by taking affine combinations of cells of the weighted Delaunay triangulation and the weighted Voronoi diagram. More precisely, for every $\mathcal{X} \subseteq \mathcal{P}$ with $\nu_{\mathcal{X}} \neq \emptyset$ a *mixed cell* is defined by

$$\mu^s_{\mathcal{X}} = s \cdot \nu_{\mathcal{X}} + (1-s) \cdot \delta_{\mathcal{X}}.$$

For $s = 0$, the mixed cell is the Delaunay cell. When s increases, it deforms affinely into the Voronoi cell for $s = 1$. We will call a mixed cell corresponding to a Delaunay ℓ-cell a mixed ℓ-cell. Since both a weighted Voronoi ℓ-cell and a weighted Delaunay $3-\ell$-cell are convex polyhedrons, it follows that a mixed ℓ-cell is also a convex polyhedron.

For $|\mathcal{X}| = 1$, the corresponding mixed cell (a mixed 0-cell) consists of all points on the corresponding Voronoi 3-cell, shrunken towards the Delaunay vertex with a factor s. A mixed 1-cell is a prism with the shrunken Voronoi facet as its base. Similarly, a mixed 2-cell is a prism with the shrunken Delaunay triangle as its base. A mixed 3-cell corresponds to a Delaunay tetrahedron shrunken towards the Voronoi vertex.

The mixed complex, associated with a shrink factor s, consists of the set of all mixed cells:

$$\text{Mix}^s\,\mathcal{P} = \{\mu^s_{\mathcal{X}} | \nu_{\mathcal{X}} \in \text{Vor } \mathcal{P}\}.$$

The mixed complex is a partition of 3-space into convex polyhedral cells.

Skin surfaces

The body of a skin surface is by definition the union of the shrunken convex hull of its weighted points $\mathcal{P}$, see Equation (2). We can extend this definition in the following way:

$$\texttt{bdy}^s\mathcal{P} \quad = \quad \bigcup\{\texttt{conv}(\mathcal{X})^s | \mathcal{X} \subseteq P\}$$

In [7] Edelsbrunner shows that only the subsets $\mathcal{X}$ for which $\nu_{\mathcal{X}} \neq \emptyset$ can touch the boundary of the skin surface. All other sets $\mathcal{X}$ with $\nu_{\mathcal{X}} = \emptyset$, are contained inside these subsets. As a result it follows that each part of the skin surface is generated by the shrunken convex hull of at most 4 points. Furthermore, he shows that a subset $\mathcal{X}$ with $\nu_{\mathcal{X}} \neq \emptyset$ generates exactly the part of the skin surface that lies inside the mixed cell $\mu^s_{\mathcal{X}}$.

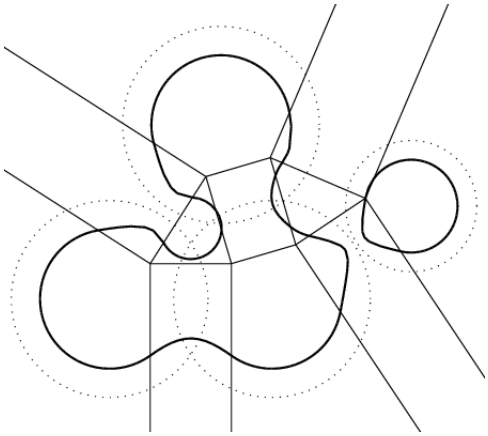

Figure 4: The skin curve of four weighted points (the dotted circles). Each mixed cell contains parts of an hyperbola or a circle

A two dimensional example is given in Figure 4. The points denote the centers of the weighted points. All rectangles (mixed cells corresponding to the Delaunay edges) contain hyperbolic patches. These hyperbolic patches bound the union of the shrunken convex hull of the two weighted points corresponding to the vertices. All other cells contain circular arcs. Depending on whether the mixed cell corresponds to a Delaunay vertex or a Delaunay face, the interior of the skin curve lies in- or outside the circle. For mixed 0-cells, the skin curve lies inside the corresponding weighted point shrunken with an factor s. On the other hand, for a mixed 2-cell, the skin curve lies outside the weighted point shrunken with a factor $1-s$, that is orthogonal to every weighted point corresponding to a vertex of the Delaunay triangle. This is the shrunken convex hull of the three weighted points corresponding to the vertices.

In $3D$, the mixed complex divides the skin surface into parts of degree two in a similar way. Assume that the affine hull through a weighted Delaunay k-cell $\delta_{\mathcal{X}}$ is aligned along the first k axes and its dual weighted Voronoi $3-k$-cell is aligned along the other axes. Then the focus lies at the origin and $\delta_{\mathcal{X}} \cap \texttt{skn}^s\mathcal{P} = \delta_{\mathcal{X}} \cap S$, where S is the implicit surface given by:

$$-\frac{1}{1-s}\sum_{i=1}^{k} x_i^2 + \frac{1}{s}\sum_{j=k+1}^{3} x_j^2 - R^2 = 0.$$

Here R^2 is the distance between the focus and a weighted point in the Delaunay cell.

For $k = 0$ and $k = 3$, the mixed cell clips a sphere, with this difference that, for $k = 0$, the sphere is convex (it contains the body), whereas for $k = 3$, it is concave. For $k = 0$, the mixed cell corresponds to a weighted Delaunay vertex. These mixed cells clip a shrunken weighted point from the input set. For $k = 1$ and $k = 2$ the mixed cell clips a hyperboloid connecting the spherical caps.

3. APPROXIMATION ALGORITHM

In this section we present the conditions imposed on the surface computed by the approximation algorithm.

The Hausdorff distance between two subsets X and Y of a Euclidean space is denoted by $d(X,Y)$. For a formal definition of the Hausdorff distance see [8].

DEFINITION 1. *Let W be a C^2-surface embedded in three space. For $r > 0$ a finite set $\mathcal{P}$ of balls with union M is r-admissible if*

(i) the boundary ∂M of M is homeomorphic to W;
(ii) every ball in $\mathcal{P}$ is M-maximal;
(iii) for each $B \in \mathcal{P}$ there is a W-maximal ball B' such that $d(B,B') \leq r$.
(iv) for each W-maximal ball $\hat{p}$ there is a $\hat{p}' \in \mathcal{P}$ such that $d(\hat{p},\hat{p}') < r$

For a given C^2-surface W the approximation algorithm takes an r-admissible set $\mathcal{P}$ of balls, and computes a C^1 *skin surface* S associated with $\mathcal{P}$ satisfying the following conditions:

$C_1(S)$: W and S are homeomorphic;

$C_2(S)$: Every ball in $\mathcal{P}$ is S-maximal;

$C_3(S)$: The Hausdorff distance $d(B_W, B_S)$ between the bodies of W and S does not exceed $c\,r$, for some positive constant c.

We realise that the conditions for an r-admissible set are rather strong, but not all restrictions are needed for the conditions to hold. In fact, we only need requirement (i) for $C_1(S)$, and (ii) for $C_2(S)$. For condition $C_3(S)$, both (iii) and (iv) are needed. Although this last condition is hard to validate, it is reasonable, since it guarantees that every ball in the Medial Axis Transform is close to a ball in the r-admissible set.

REMARK 2. The construction of the admissible set $\mathcal{P}$ of W-maximal balls is difficult, and to the best of our knowledge there is no general method yet. However, we conjecture that the *power crust algorithm* [1] can be adapted to construct a set of r-admissible balls from a suitable sample of the surface W. The latter algorithm constructs a polyhedral surface with the same topology as the surface W, provided the sample is sufficiently dense. More precisely, recall that the local feature size (LFS) at a point $x \in W$ is the distance of x to the medial axis of W. A set of sample points is a ϱ-sample if for every point $x \in W$ there is a sample point at distance at most $\varrho \cdot LFS(x)$ from x. For a ϱ-sample with $\varrho \leq 0.1$, the power crust algorithm yields a set of approximate maximal balls for which the union of the boundary is homeomorphic to W.

DEFINITION 3. *For a set of balls $\mathcal{P}$ the surface $S_s(\mathcal{P})$, $0 < s \leq 1$, is the skin surface with shrink factor s associated with the set of balls $\mathcal{P}^{1/s}$ defined by*

$$\mathcal{P}^{1/s} = \{\hat{p}^{1/s} | \hat{p} \in \mathcal{P}\}.$$

If $\mathcal{P}$ is an r-admissible set of balls then conditions $C_1(S_1)$ and $C_2(S_1)$ are satisfied. Indeed, $\mathcal{P}^1 = \mathcal{P}$, so S_1 is the boundary of the union of these balls. Therefore, conditions $C_1(S_1)$ and $C_2(S_1)$ are trivially satisfied. In Section 6, we show that for the region B_W enclosed by W we have $d(B_W, \cup\mathcal{P}) \leq r$, and it follows that $C_3(S_s)$ is also satisfied for $s = 1$.

Obviously, S_1 is *not smooth*. However, for values of s slightly smaller than 1 conditions $C_1(S_s)$ and $C_2(S_s)$ are still satisfied, and S_s is a C^1-surface. Indeed, for s near 1, the surface S_s is homeomorphic to S_1, and hence to W. Since $(\hat{p}^{1/s})^s = \hat{p}$, it follows that $\mathcal{P}$ is a collection of S_s-maximal balls, for s slightly smaller than 1.

One or more of these conditions may be violated for smaller values of s. See Figure 5 for a sequence of skin curves corresponding to decreasing values of s. Figure 5a shows the skin curve for $s = 1$. As we decrease the shrink factor, the skin curve becomes tangent continuous, due to the appearance of hyperbolic patches connecting the circles, see Figure 5b.

As these patches grow further, they eventually cause a change in the topology of the skin curve, see Figure 5d, i.e., a violation of condition $C_1(S_s)$. Finally, as we decrease the shrink factor even further, the balls are no longer maximal, as depicted in Figure 5e. In Section 5 we derive, how far the shrink factor can be decreased without occurrence of this phenomenon.

Therefore, our goal is to determine the interval of s-values for which conditions $C_1(S_s)$–$C_3(S_s)$ are satisfied. One of the main results of this paper is

THEOREM 1. *The value*

$$s_i(\mathcal{P}) = \inf\{s \mid 0 \leq s < 1 \text{ and condition } C_i(S_s) \text{ is satisfied}\}$$

can be computed in $O(N^2 \log N)$ time and $O(N^2)$ space, where N is the number of balls in $\mathcal{P}$.

The algorithm presented in Section 4 computes $s_1(P)$ by maintaining the weighted Voronoi diagram of a set of balls with growing radius; This algorithm is of some independent interest. In Section 5 this method is slightly adapted to compute an interval of s-values for which condition $C_2(S_s)$ holds. Finally, in Section 6, we determine the error with respect to the Hausdorff distance between the input surface and the computed skin surface. This error analysis gives us a value for $s_3(\mathcal{P})$.

4. MAINTAINING THE TOPOLOGY OF THE UNION OF GROWING BALLS

In this section we discuss an algorithm to compute the smallest τ such that the boundary of the union of the balls $\mathcal{P}^\tau$ is not homeomorphic to $\partial\bigcup\mathcal{P}$. Since the skin surface $S_s(\mathcal{P})$ is homeomorphic to the boundary of the union of the set of balls $\mathcal{P}^{1/s}$, cf Section 2, this is equivalent to computing the value $s_1(\mathcal{P})$, defined in Section 3. Figure 6 illustrates this equivalence in 2D, where the change of topology of the skin curve coincides with the change of topology of the balls defining the skin curve.

Let $\mathcal{Q}$ be a set of balls in $\mathbb{R}^3$, and, for $t \in \mathbb{R}$, let $\partial(\cup\mathcal{Q}^t)$ denote the boundary of the union $\cup\mathcal{Q}^t$.

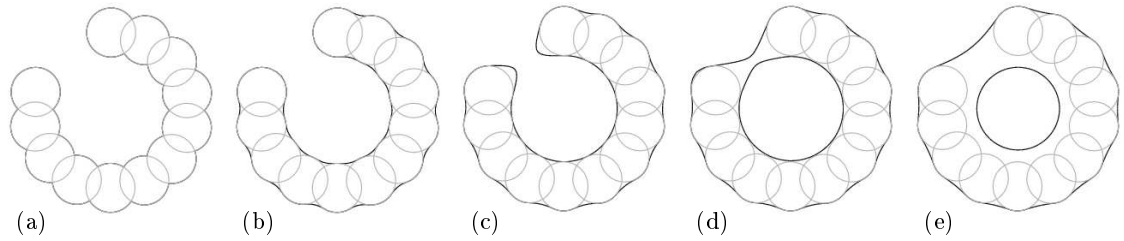

Figure 5: The set of maximal circles and the approximating skin curve for different shrink factors (decreasing from left to right).

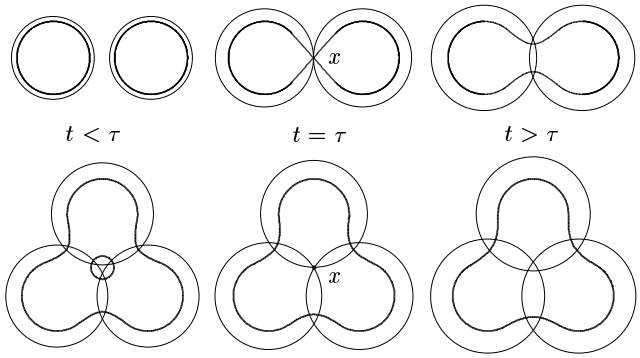

Figure 6: Change of topology at (x, τ) of a union of two (top row) and three (bottom row) weigthed points.

DEFINITION 4. *$\partial(\cup\mathcal{Q}^t)$ **changes topology** at $(x, \tau) \in \mathbb{R}^3 \times \mathbb{R}$ if for every neighborhood U of x in $\mathbb{R}^3$ there is an $\varepsilon > 0$ such that, for $\tau - \varepsilon < t_1 < \tau < t_2 < \tau + \varepsilon$, the sets $U \cap \partial(\cup\mathcal{Q}^{t_1})$ and $U \cap \partial(\cup\mathcal{Q}^{t_2})$ are not homeomorphic.*
*A ball $\hat{q} \in \mathcal{Q}$ is **involved** in the change of topology at (x, τ) if x is a point of the ball $\hat{q}^\tau$.*

Obviously, if $\partial(\cup\mathcal{Q}^t)$ changes topology at (x, τ), then $x \in \partial(\cup\mathcal{Q}^\tau)$. Occasionally, we just say that $\partial(\cup\mathcal{Q}^t)$ changes topology for $t = \tau$ if the location x is irrelevant.

LEMMA 2. *Suppose all weighted points of $\mathcal{Q}$ are involved in a change of topology of $\partial(\cup\mathcal{Q}^t)$ at (x, τ). Then*

$$\cap\mathcal{Q}^t = \begin{cases} \emptyset, & \text{for } t < \tau, \\ \{x\}, & \text{for } t = \tau, \end{cases}$$

and $\cap\mathcal{Q}^t \neq \emptyset$, for $t \geq \tau$.

PROOF. Suppose the interiors of the balls in $\mathcal{Q}^\tau$ intersect in a point y. Then y belongs to all balls in $\mathcal{Q}^{t'}$, for some $t' < \tau$. Since the weights of all weighted points in $\mathcal{Q}^t$ are increasing with t, whereas the centers are independent of t, it follows that, for $t > t'$, $\cup\mathcal{Q}^t$ is starshaped with respect to y. In particular, for $t > t'$, the set $\partial(\cup\mathcal{Q}^t)$ is non-empty, and does not change topology. This contradiction shows that the interiors of the balls in $\mathcal{Q}^\tau$ are disjoint, and hence $\cap\mathcal{Q}^\tau$ consists of a single point, which, by definition, is equal to x. For $t < \tau$ the ball $\hat{q}^t$ is contained in the interior of $\hat{q}^\tau$, so $\cap\mathcal{Q}^t = \emptyset$. Since x is contained in the interior of all balls $\hat{q}^t$ for $t > \tau$, it finally follows that $\cap\mathcal{Q}^t \neq \emptyset$, for $t \geq \tau$. □

LEMMA 3. *Suppose $\partial(\cup\mathcal{P}^t)$ changes topology at $(x, \tau) \in \mathbb{R}^3 \times \mathbb{R}$. Then the subset $\mathcal{Q}$ of $\mathcal{P}$ of weighted points involved in this change of topology defines a cell in the weighted Voronoi diagram of $\mathcal{P}^\tau$. This cell contains the point x.*

PROOF. It follows from Lemma 2 that x belongs to the boundary of the weighted points in $\mathcal{Q}$, so $\pi(\hat{q}^\tau, x) = 0$, for $\hat{q} \in \mathcal{Q}$. By definition, $\pi(\hat{p}^\tau, x) > 0$, for $\hat{p} \in \mathcal{P} \setminus \mathcal{Q}$, so x belongs to the cell $\nu_{\mathcal{Q}^\tau}$ of the weighted Voronoi Diagram of $\mathcal{P}^\tau$. □

Consider a set $\mathcal{Q}$ of balls, with boundaries intersecting at x. Generically, the tangent planes at x intersect transversally, i.e., the intersection is an affine space of codimension $|\mathcal{Q}|$. The balls $\mathcal{Q}^t$ form a one-parameter family, so at isolated values of t we expect intersections with non-transversal tangent planes, i.e., the codimension of this intersection *is less than* $|\mathcal{Q}|$. We impose the following *generic condition* on the family $\mathcal{P}$.

Generic change of topology. *Suppose $\partial(\cup\mathcal{P}^t)$ changes topology at (x, τ), and the subset $\mathcal{Q}$ of $\mathcal{P}$ consists of the balls involved in this change of topology. Then the tangent planes at x of the balls in $\mathcal{Q}$ intersect in an affine space of codimension $|\mathcal{Q}| - 1$.*

In the plane, at most three weighted points are involved in a generic change of topology, see Figure 6. When two weighted points are involved, the change in topology corresponds to the *creation of a bridge* between two parts of the boundary. A change of topology in which three weighted points are involved corresponds to the *filling of a void*.

Let $\mathcal{Q}$ be the set of balls involved in a generic change of topology in three-space, then $|\mathcal{Q}| \leq 4$. These changes of topology correspond to the creation of a bridge, when $|\mathcal{Q}| = 2$, the filling of a tunnel, when $|\mathcal{Q}| = 3$, or the filling of a void, when $|\mathcal{Q}| = 4$.

Lifting the weighted Voronoi Diagram.

In view of Lemma 3 changes of topology of $\partial(\mathcal{P}^t)$ are related to the cells of the weighted Voronoi Diagram. To incorporate t-dependence, we lift the weighted Voronoi Diagram to $\mathbb{R}^3 \times \mathbb{R}$. To this end consider, for $\hat{p}, \hat{q} \in \mathcal{P}$, the half space $H^*_{\hat{p},\hat{q}} \subset \mathbb{R}^3 \times \mathbb{R}$ defined by

$$\begin{aligned} H^*_{\hat{p},\hat{q}} &= \{(x,t) \in \mathbb{R}^3 \times \mathbb{R} \mid \pi_{\hat{p}^t}(x) \leq \pi_{\hat{q}^t}(x)\} \\ &= \{(x,t) \in \mathbb{R}^3 \times \mathbb{R} \mid \\ &\qquad \langle x, q - p\rangle + \tfrac{1}{2}t(\mathrm{Q} - \mathrm{P}) \leq \tfrac{1}{2}\| q \|^2 - \tfrac{1}{2}\| p \|^2\}. \end{aligned}$$

The boundary of this half space is called the *extended bisector* of the weighted points $\hat{p}$ and $\hat{q}$. With a weighted point $\hat{p} \in \mathcal{P}$ we associate the extended Voronoi cell $V^*(\hat{p})$ in $\mathbb{R}^3 \times \mathbb{R}$ defined by

$$V^*(\hat{p}) = \bigcap_{\hat{q} \,:\, \hat{q} \neq \hat{p}} H^*_{\hat{p},\hat{q}}.$$

For $\hat{p} = (p, \mathrm{P}) \in \mathcal{P}$, the point $(p, 0)$ belongs to $V^*(\hat{p})$, so no extended Voronoi cell is empty. The extended Voronoi cells determine a subdivision of $\mathbb{R}^3 \times \mathbb{R}$ into convex polyhedra. This subdivision is called the *extended Voronoi Diagram* of $\mathcal{P}$, and is denoted by $VD^*(\mathcal{P})$. The weighted Voronoi cell of $\hat{p}^\tau$ with respect to $\mathcal{P}^\tau$ is the intersection of the extended Voronoi cell $V^*(\hat{p})$ and the hyperplane $t = \tau$.

General position. *We assume that the weighted points in $\mathcal{P}$ are in general position in the sense that*
(i) no $k+2$ weighted points lie on a common k-flat, for $0 \leq k < 3$ (this is equivalent to the centers of the weighted points being in general position);
(ii) no 6-tuple of weighted points has a common ortosphere;
(iii) a 5-tuple of weighted points has a common ortosphere for a finite set of values of t.

If $\mathcal{P}$ is in general position, then, for $0 \leq k \leq 4$, a k-face of the extended Voronoi Diagram is defined by $5-k$ weighted points, in other words, it is incident to $5-k$ extended Voronoi cells (a cell is a 4-face of the extended Voronoi Diagram). Since a vertex is defined by 5 weighted points, it is incident upon $\binom{5}{4} = 5$ edges. If the number of edges incident upon a vertex (x, τ) of the extended Voronoi Diagram, and contained in the half space below the hyperplane $t = \tau$, is $k+1$, $0 \leq k \leq 3$, then this vertex is of *type k*. Since the radius of the balls in $\mathcal{P}^t$ is increasing with t starting at $t = 0$, we do not encounter vertices of type 0.

LEMMA 4. *The extended Voronoi Diagram of a set of N weighted points in $\mathbb{R}^3$ has $O(N^2)$ faces.*

PROOF. Obviously, the number of 4-faces is $O(N)$, since there is a one-one correspondence between the set of 4-faces and the set of weighted points. Similarly, the number of 3-faces is $O(N^2)$, since each 3-face corresponds to a pair of weighted points, and distinct 3-faces correspond to distinct pairs.

Vertices of type 0 (type 3) are t-minimal (t-maximal) points of a 4-face, so there are $O(N)$ vertices of type 0 and of type 3. Vertices of type 1 (type 2) are t-minimal (t-maximal) points of a 3-face, so there are $O(N^2)$ vertices of type 1 and of type 2.

Since generically there is a constant number of 1-faces and 2-faces incident upon a vertex, the total number of faces of the extended Voronoi Diagram is $O(N^2)$. □

To see that the quadratic complexity is worst case optimal, we construct a set of weighted points in $2D$ for which the structure of the weighted Voronoi diagram changes a quadratic number of times as we grow the weighted points. More specifically, $O(N)$ cells get an empty Voronoi cell (corresponding to vertices of type 3 in the extended $3D$ Voronoi diagram), and $O(N^2)$ edge flips (these are vertices of type 1 or 2) occur.

To construct this set, align the first $N/2$ points along the vertical axis with a constant weight P. Align the other weighted points along the horizontal axis such that P′, with $\mathrm{P}' < \mathrm{P}$. Further ensure that the Voronoi cells of two of these weighted points touch the Voronoi cells of all weighted points aligned along the vertical axis. As a result the Voronoi cells of all other weighted points touch only the two Voronoi cells of the outer weighted points that are aligned vertically. This set of weighted points is depicted in Figure 7(a), and its Voronoi diagram in Figure 7(b).

Since the weights of the vertically aligned weighted points are equal, a common Voronoi edge between two of these points does not change as we grow the weighted points. The same holds for an edge generated by two weighted points on the horizonal axis. Because of the difference in weight of points on the horizontal and vertical axis, the Voronoi cells of the weighted points on the vertical axis encroach upon the the Voronoi cells of the weighted points on the horizontal axis. As a result for each weighted point on the horizontal axis, we have $O(N)$ edge flips before the Voronoi cell becomes empty. The two outmost weighted points form an exception, since their Voronoi cells do not become empty. This process is depicted in Figure 7(b) to (d).

PROPOSITION 5. *The extended Voronoi Diagram of a set of N weighted points in $\mathbb{R}^3$ can be constructed in $O(N^2 \log N)$ time.*

PROOF. We maintain the weighted Voronoi Diagram of $\mathcal{P}^t$ using a sweep-hyperplane algorithm, starting at $t = 0$. The weighted Voronoi Diagram for $t = 0$ can be computed in $O(N^2)$ time, using the algorithm from [9]. The combinatorial structure of $VD(\mathcal{P}^t)$ changes at the t-coordinates of the vertices of the extended Voronoi Diagram. There are several types of *events*, depending on the type of the vertex: the type of the event is the type of the vertex. At an *event of type k*, $0 < k \leq 3$, a k-face f (and its incident j-faces, $0 \leq j < k$) is destroyed, and a $3-k$-face f' is created.

Since the weighted points are in general position, face f is a k-simplex just before its destruction. Therefore, the algorithm maintains a priority queue storing the destruction times of all k-simplices in the current weighted Voronoi Diagram. A Voronoi vertex corresponds to a line segment in the extended Voronoi diagram. By definition of the extended Voronoi diagram, a Voronoi vertex parametrized in terms of t, moves along a line in $\mathbb{R}^3$ at constant speed. At an event, at least two Voronoi vertices coincide. It follows that computing the destruction time of a Voronoi cell is equivalent to computing the time at which two of its vertices coincide.

Processing an event of type k now boils down to the following update operations:
(i) Update the weighted Voronoi Diagram, i.e., remove the destroyed k-simplex and insert the created $3-k$-simplex, thereby updating all incidence relations;

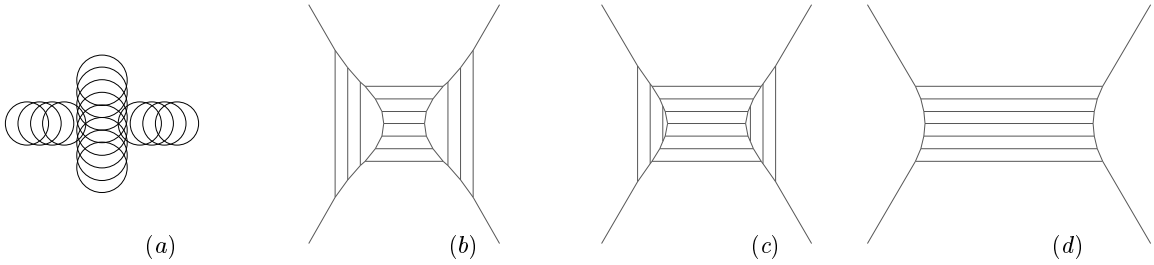

Figure 7: A set of weighted points for which the Voronoi diagram changes topology a quadratic number of times. Figure (a) shows the set of weighted points and (b) its Voronoi diagram. Figure (c) and (d) show the Voronoi diagram for increasing t.

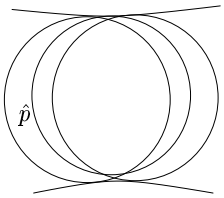

Figure 8: Three weighted points in an r-admissible set $\mathcal{P}$. The radius of the maximal circle $\hat{p}$ in the middle is shrunken with a factor r. The result is that $\hat{p}$ only touches $\bigcap \mathcal{P}$ in one small circular arc.

(ii) Update the priority queue by adjusting the destruction times of all j-simplices, $j > 0$, incident with the newly created simplex.

The destruction time of a k-simplex is equal to the destruction time of its edges. Therefore, it is sufficient to compute the destruction time of each new or updated edge. Since only a constant number of faces is involved in a single event, step (i) takes $O(1)$ time, and of step (ii) takes $O(\log N)$ time.

Since multiple faces are incident upon the destroyed and created simplices, time stamps in the priority queue correspond to multiple simplices. However, in view of Lemma 4 the total number of time stamps is $O(N^2)$. Therefore the time complexity of the algorithm is $O(N^2 \log N)$. □

COROLLARY 6. *Let $\mathcal{P}$ be a set of N weighted points in $\mathbb{R}^3$. Then τ_0, defined by*

$$\tau_0 = \min\{\tau \mid \partial(\mathcal{P}^\tau) \text{ changes topology for } t = \tau\},$$

can be computed in $O(N^2 \log N)$ time.

PROOF. Consider a face of the extended Voronoi Diagram, defined by a subset $\mathcal{Q}$ of $\mathcal{P}$. There is a unique value t for which $\cap\mathcal{Q}^t$ consists of a single point, x say: solve (x, t) from the equation $\pi(\hat{q}^t, x) = 0$, for some, and hence all, $\hat{q} \in \mathcal{Q}$. In this way we associate a unique time stamp $\tau_\mathcal{Q}$ with each face of the extended Voronoi Diagram. The value τ_0 is the minimal time stamp greater than 1. Since there are $O(N^2)$ faces, it can be computed from the extended Voronoi Diagram in $O(N^2)$ additional time. □

5. PRESERVING MAXIMALITY OF THE BALLS

The medial axis transform forms a good shape descriptor for surfaces. From a surface W we obtained an r-admissible set $\mathcal{P}$ of approximate maximal balls of W. In this section, we consider this set to be a shape descriptor for W. We approximate W with a surface that is a good approximation in the sense that each maximal ball in $\mathcal{P}$ is also maximal in the approximating surface.

Before we construct such a surface, first we give a formal definition of a T-maximal ball w.r.t. some surface T.

DEFINITION 7. *A ball is T-maximal if it is enclosed in the surface T and it is not contained in any other ball enclosed in T.*

Consider an r-admissible set $\mathcal{P}$. The surface $M = \bigcup \mathcal{P}$ is composed of spherical patches. Each spherical patch γ is part of the boundary of some maximal ball in $\mathcal{P}$.

Requirement (ii) of an r-admissible set ensures that each weighted point generates at least one spherical patch on $\bigcup \mathcal{P}$. In Figure 8, we show that for any $r > 0$, we can construct an r-admissible set that contains a weighted point $\hat{p}$ such that $\hat{p}$ generates only one circular arc and this arc is arbitrary small. This example can be extended in any dimension. Therefore to guarantee that each weighted point in $\mathcal{P}$ is also maximal in the approximating skin surface, we have to guarantee that for every spherical patch γ, the approximating skin surface contains a subset of γ. As long as this subset exists, it ensures that its generating weighted point is maximal.

Conceptually, we start with the skin surface for $s = 1$. For this approximation, we know that $\gamma \cap S_1 = \gamma$. As we decrease s, patches between spheres arise and $\gamma \cap S_s$ becomes a proper subset of γ. While doing so, the intersection of γ and the skin surface can become disconnected. When this happens, we consider each of these components seperately. We stop decreasing s just before some γ touches S_s only in one point.

To determine when a patch γ only touches the skin surface S_s in a point, we use the mixed complex. Recall that the mixed complex decomposes the skin surface into quadratic patches. In particular, a mixed 0-cell clips its corresponding maximal ball. A mixed 0-cell corresponds to a weighted

point $\hat{p} \in \mathcal{P}$. Its shape is the Voronoi 3-cell corresponding to $\hat{p}$ shrunken with a factor s towards the weighted Delaunay 0-cell $\delta_{\{\hat{p}\}}$.

LEMMA 5. *Let γ be a spherical patch on $\bigcup \mathcal{P}$ generated by a weighted point $\hat{p} \in \mathcal{P}$. Then there is a weighted Delaunay 3-cell $\delta_{\mathcal{X}}$ adjacent to $\delta_{\{\hat{p}\}}$ for which the segment between p and $\nu_{\mathcal{X}}$ intersects γ. Further, the skin surface S_s contains a spherical patch of γ iff*

$$s \cdot |p - \nu_{\mathcal{X}}| > \text{P}$$

PROOF. First, we analyse the shape of a mixed 0-cell more carefully. A Voronoi 3-cell is a convex polyhedron with its Voronoi 0-cells at its vertices. Since a mixed 0-cell is a Voronoi 3-cell shrunken towards the center of its corresponding weighted point, a mixed 0-cell is a convex polyhedron with its vertices on the line from the weighted Delaunay 0-cell to each Voronoi 0-cell.

To be more precise, for every weighted Delaunay 3-cell $\delta_{\mathcal{X}}$ adjacent to $\delta_{\{\hat{p}\}}$, the mixed 0-cell has a vertex at the point:

$$s \cdot \nu_{\mathcal{X}} + (1 - s) \cdot \delta_{\{p\}}.$$

The intersection of the spherical cap γ with S_s is non-empty if some part of γ lies inside the mixed 0-cell. Since a mixed 0-cell is a convex polyhedron, this is the case if a vertex of the polyhedron lies outside $\hat{p}$. Therefore the intersection of γ with the skin surface is non-empty if the inequality in the lemma holds. Note that in this equation $\nu_{\mathcal{X}}$ also depends on the shrink factor. □

The algorithm to obtain the lowest shrink factor is similar to the algorithm used to compute the minimal shrink factor for which the skin surface is homeomorphic to the original surface W. Instead of testing when a set of weighted points in a weighted Delaunay cell causes a topological change, we test for each weighted Delaunay vertex when a spherical patch degenerates into a point. Just before this happens, we have obtained the lowest shrink factor for which we can guarantee that the maximal balls in $\mathcal{P}$ are also maximal in the skin surface.

6. ERROR ESTIMATES

In this section we construct a minimal shrink factor s_3 such that, given an r-admissible sample of a surface W, the Hausdorff distance between the body B_W bounding W and the body $B_S(s)$ of the approximating skin surface S_s is at most $2r$ for $s_3 \leq s \leq 1$. We do this in the following way: first, we show that under the conditions in Definition 1 (in particular condition (iv)), the distance between B_W and $\cup\mathcal{P}$ is at most r. Then we construct a skin surface S_s for which we guarantee that the distance from each patch to $\bigcup \mathcal{P}$ is at most r. It follows from these results that the distance between B_W and the body of the approximated skin surface is at most $2r$.

LEMMA 6. *Let W be a closed surface bounding a body B_W. For an r-admissible sample $\mathcal{P}$ of W, we have that:*

$$d(B_W, \bigcup \mathcal{P}) \leq r,$$

PROOF. Let $x \in B_W$, and $\hat{p}$ be an W-maximal ball containing x. Using restriction (iv), we know that there is a ball $\hat{p}'$ in $\mathcal{P}$ with $d(\hat{p}, \hat{p}') \leq r$. It immediately follows that $d(x, \hat{p}') \leq r$ and therefore $d(x, \bigcup \mathcal{P}) \leq r$.
Using condition (iii) and a similar reasoning, it follows that a point $x \in \bigcup \mathcal{P}$ lies at most at a distance r from B_W. □

It follows from this lemma that S_1 satisfies $C_3(S)$. In the rest of this section we determine the minimal shrink factor s_3 for which we can show that the maximal distance between the body $B_S(s)$ of S_s and $\bigcup \mathcal{P}$ is at most r. Since $\bigcup \mathcal{P} \subseteq B_S(s)$, it follows that $d(x, B_S(s)) = 0$ for all $x \in \bigcup \mathcal{P}$. To show the converse, that $d(x, \bigcup \mathcal{P}) \leq r$ for all $x \in B_S(s)$, we use the decomposition of S_s by the mixed complex.

The mixed complex decomposes a skin surface into quadratic patches. Each mixed cell $\mu^s_{\mathcal{X}}$ contains the quadratic patch that is the intersection of $\mu^s_{\mathcal{X}}$ with $\bigcup \texttt{conv}(\mathcal{X})^s$. A set $\mathcal{X}$ can only generate a part of the skin surface if it generates a non-empty weighted Voronoi Cell.

Instead of taking the convex hull of a set of weighted points $\mathcal{X}$, consider the affine hull of $\mathcal{X}$. Let $\hat{f}$ be the weighted point in the affine hull with minimal weight. The weight of $\hat{f}$ is positive if and only if its center is contained in $\bigcup \mathcal{X}$. Simple calculations show that f is the focus of $\mathcal{X}$.

We devide the sets $\mathcal{X}$ that can generate parts of the skin surface in two classes. First, we consider the sets for which $\hat{f}$ has positive weight. For these sets we bound the distance from the part of the skin surface generated by $\mathcal{X}^{1/s}$ to $\bigcup \mathcal{X}$ by r. We can not do this if the weight of $\hat{f}$ is negative. Therefore we make sure that these cells do not contribute to any part of the skin surface.

Bounding the distance of a patch to $\bigcup \mathcal{P}$

Let $\mathcal{X} \subseteq \mathcal{P}$ form a non-empty weighted Voronoi cell, and let $\hat{f} \in \texttt{aff}(\mathcal{X})$ be the weighted point with its center at the focus of $\mathcal{X}$. Assume that $\hat{f}$ has positive weight. It follows that each point in the affine hull, and therefore also in the convex hull, of $\mathcal{X}$ has positive weight.

Now let $X = \{p \mid \hat{p} \in \mathcal{X}\}$ and let $m \in \texttt{aff}(X)$ be the point equidistant to each $x \in X$.

LEMMA 7. *If the weight of $\hat{f}$ is positive, then for each shrink factor $s_{\mathcal{X}} \leq s \leq 1$ with*

$$s_{\mathcal{X}} = 1 - \frac{r^2}{|m - x|^2}, \tag{4}$$

for some $x \in X$, we have that

$$d(\bigcup \mathcal{X}, S_s(\mathcal{X})) \leq r.$$

PROOF. We define the function $\text{P}_s(p)$ as the weight of the weighted point in $\texttt{conv}(\mathcal{X}^{1/s})^s$ centered at p.

The set $\Pi(\texttt{conv}(\mathcal{X}))$ forms a hyperplane in $\mathbb{R}^4$. Since the weight is defined as the distance along the last coordinate

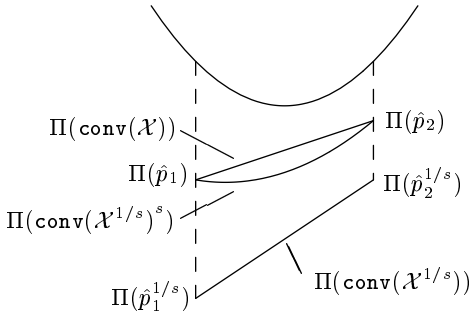

Figure 9: The vector space obtained under the lifting operator Π on 1-dimensional weighted points. It shows how the convex hull of the set $\mathcal{X} = \{\hat{p}_1, \hat{p}_2\}$ deforms as we decrease the shrink factor.

axis to the unit paraboloid, $\mathrm{P}_1(p)$ is a paraboloid with leading coefficient 1. For the analogue of the vector space obtained under Π of weighted points with an 1-dimensional center, see Figure 9. The top parabola denotes the set of weighted points with zero weight.

From a similar reasoning it follows that $\Pi(\mathtt{conv}(\mathcal{X}^{1/s}))$ also a hyperplane forms in $\mathbb{R}^4$. By the definition of shrinking, $\mathrm{P}_s(p)$ is a paraboloid with leading coefficient s.

Knowing the leading coefficient of $\mathrm{P}_s(p)$, the paraboloid is fixed upto translation. Therefore consider a weighted point $\hat{p'} \in \mathcal{X}$. It is first shrunken with a factor $1/s$ and then with a factor s, therefore $\mathrm{P}_s(p') = \mathrm{P}'$, as is also depicted in Figure 9. Using these points the parabola is uniquely defined, and it follows that

$$\mathrm{P}_s(p) = \mathrm{P}_1(p) - (1-s)(|x-m|^2 - |p-m|^2),$$

for some $x \in X$, since $|x, m| = |x', m|$ for $x, x' \in X$. Indeed, $\mathrm{P}_s(p)$ is a paraboloid with leading coefficient s, since $\mathrm{P}_1(p)$ has leading coefficient 1. Further, the weight of the points in $\mathcal{X}$ is independent of the shrink factor.

Using the fact that for $s = 1$, $\hat{p}(1) \in \mathtt{conv}(\mathcal{X})$ is contained in $\bigcup \mathcal{X}$ and that each weighted point corresponds to a ball with real radius, we bound the patch in terms of the distance between $\hat{p}(s) \in \mathtt{conv}(\mathcal{X}^{1/s})^s$ and $\hat{p}(1)$.

$$\begin{aligned} d(\hat{p}(s), \bigcup \mathcal{P}) &\leq d(\hat{p}(s), \hat{p}(1)) \\ &= \sqrt{\mathrm{P}(s)} - \sqrt{\mathrm{P}(1)} \\ &\leq \sqrt{(1-s)(|x-m|^2 - |p-m|^2)} \end{aligned}$$

The equation in the lemma follows if we take $p = m$ to obtain the maximal difference in weight, and then solve $d(\hat{p}(s), \bigcup \mathcal{P}) = r$ for s. Note that this bound can be too pessimistic since m can lie outside the convex hull of X, and therefore $|p-m|^2 > 0$ for all $p \in \mathtt{conv}(X)$ □

Avoiding caps in a skin surface

In the case that we have a set $\mathcal{X} \in \mathcal{P}$ for which the weight of $\hat{f}$ is negative, the previous lemma does not hold. To see this, consider a weighted points $\hat{p}$ in the convex hull of $\mathcal{X}$ with negative weight. The center of his point lies outside $\bigcup \mathcal{X}$. As we decrease the shrink factor the weight may become real, but we can not bound the distance from this maximal ball to $\bigcup \mathcal{X}$ in a similar way.

In order to avoid this situation to happen we assure that the mixed cell corresponding to $\mathcal{X}$ does not contain any part of the skin surface. For a mixed 0-cell corresponding to a weighted point $\hat{p}$, the focus is located at the center of $\hat{p}$ and has therefore always a positive weight.

Consider a mixed 1-cell corresponding to a set $\mathcal{X}$. The patch generated by this cell is a two sheeted hyperboloid. Its symmetry axis is the line through the centers of the two weighted points in $\mathcal{X}$. For $s = 1$, we know that the mixed 1-cell does not contain any part of the skin surface, since no patches are generated. As we decrease the shrink factor, the first point where the skin surface intersects the mixed 1-cell is on the symmetry axis.

A similar analysis of the mixed 2-cells and 3-cells, shows that the first point where the skin surface will intersect the mixed cell lies on the segment between the focus f of the corresponding set of weighted points $\mathcal{X}$ and the center of a weighted point in $\mathcal{X}$. Moreover, the skin surface intersects the mixed cell if

$$s \cdot |p - f| \leq \mathrm{P}, \tag{5}$$

for some $\hat{p} \in \mathcal{X}$. This can be shown in a similar way as done in the proof of Lemma 5 and is left as an exercise to the reader.

Summarising, in Lemma 6 we showed that the distance between an arbitrary surface W and the union of an r-admissible sample $\mathcal{P}$ smaller is than r. In the rest of the section we obtained two equations.
The first equation, Equation (4), holds if $\hat{f}$ has positive weight and gives us a minimal shrink factor for which we can prove that the distance from the body of $S_s(\mathcal{X})$ to $\bigcup \mathcal{X}$ is at most r. The second equation, Equation (5), gives us the minimal shrink factor for which $\mathcal{X}$ does not generate any part of $S_s(\mathcal{P})$, if the weight of $\hat{f}$ is negative.

Combining these two results, we obtain a value s_3 for which the Hausdorff distance between the body bounded by W and the body of the skin surface S_s, with $s_3 \leq s \leq 1$, is at most $2r$.

7. CONCLUSION AND FUTURE WORK

We presented an algorithm to effectively compute a C^1-approximation S of a surface W represented by a set of approximate W-maximal balls. The approximation S is a skin surface, which is homeomorphic to the boundary of the union of the approximate W-maximal balls. Furthermore, the Hausdorff distance between the regions enclosed by W and S converges to zero as we increase the density of the sample of maximal balls.

A disadvantage of our method is that the surface is usually bumpy, i.e., the error of the tangent vector (the C^1-error) is not bounded, since we will always have concave patches inbetween two balls. Another drawback of our algorithm is that it determines the shrink factor globally: if a high shrink factor is needed at one part of the skin surface, this influences the approximation of the whole surface.

For s close to 1, the skin surface and the boundary of the union of the input balls is almost the same. This would imply that our approach does not improve on the union of the balls. We assume that the shrink factor will be significant smaller than 1. In fact, we conjecture:

CONJECTURE 8. *For a C^2-surface W there is a $s_W < 1$ such that the shrink factor s_S corresponding to a sample of the Medial Axis Transform of W converges to s_W if the sampling density goes to* 1.

We think that the value s_W depends upon the maximal ratio between the radius of the maximal ball touching a point $x \in W$ and the local feature size in x. The larger this ratio, the higher s_W. We are currently trying to prove this conjecture.

We are currently investigating adaptations of the methods presented in this paper, yielding C^1-surfaces that are less bumpy. We further plan to improve our algorithm such that it chooses the shrink factor adaptively, i.e., based on local constraints. Another topic of future research is the construction of an r-admissible set of (approximate) W-maximal balls.

8. REFERENCES

[1] N. Amenta, S. Choi, and R. Kolluri. The power crust, unions of balls, and the medial axis transform. *Internat. J. Comput. Geom. Appl.*, 19:127–153, 2001.

[2] J.-D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp. Geometry Theory and Applications*, 2003. to appear.

[3] J.-D. Boissonnat, O. Devillers, and J. Duquesne. Computing Connolly Surfaces. In *IFIP Conference on Algorithms and efficient computation*, 1992.

[4] H. Cheng, T. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete & Computational Geometry*, 25:525–568, 2001.

[5] H.-L. Cheng, H. Edelsbrunner, and P. Fu. Shape space from deformation. *Comput. Geom. Theory Appl.*, 19:191–204, 2001.

[6] S. Cheng, H. Edelsbrunner, P. Fu, and K. Lam. Design and analysis of planar shape deformation. *Proc. 14th Ann. Sympos. Comput. Geom.*, 1998.

[7] H. Edelsbrunner. Deformable smooth surface design. *Discrete & Computational Geometry*, 21:97–115, 1999.

[8] J. E. Goodman and J. O'Rourke, editors. *Handbook of Discrete and Computational Geometry*. CRC Press LLC, Boca Raton, FL, 1997.

[9] B. Joe. Construction of three-dimensional Delaunay triangulations using local transformations. *Comput. Aided Geom. Design*, 8(2):123–142, May 1991.