University of Groningen

Approximation by skin surfaces

Kruithof, Nico; Vegter, Gert

Published in:
Computer-Aided design

DOI:
[10.1016/j.cad.2004.01.007](https://doi.org/10.1016/j.cad.2004.01.007)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2004

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Kruithof, N., & Vegter, G. (2004). Approximation by skin surfaces. *Computer-Aided design*, 36(11), 1075-1088. <https://doi.org/10.1016/j.cad.2004.01.007>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

ELSEVIER

Computer-Aided Design 36 (2004) 1075–1088

COMPUTER-AIDED
DESIGN

www.elsevier.com/locate/cad

Approximation by skin surfaces

Nico Kruithof*, Gert Vegter

Department of Mathematics and Computing Science, University of Groningen, Groningen, The Netherlands

Accepted 9 January 2004

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 patches of spheres and hyperboloids 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.

© 2004 Elsevier Ltd. All rights reserved.

Keywords: Surface approximation; Skin surfaces; Medial axis; Maximal balls; Power diagram; Regular triangulation; Error bound

1. Introduction

We consider the problem of approximating a simple, regular smooth (C^2) closed surface in $\mathbb{R}^3$ by a skin surface. Skin surfaces, introduced by Edelsbrunner in Ref. [13], 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 parameterized 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 patches of spheres and hyperboloids connecting the balls. An example in 2D of a curve, reconstructed using this approach, is drawn in Fig. 1c. In Fig. 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 a surface W is a finite sample of the medial axis transform of W , i.e. a finite set of W -maximal balls. Let Σ be a closed surface in $\mathbb{R}^3$ (compact without boundary), which is the boundary of a compact subset V of $\mathbb{R}^3$. A ball is called an *empty ball* of Σ if it is contained in the closure of V . A Σ -*maximal ball* is an empty ball of Σ not contained in any other empty ball of Σ . 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 strong deformation retract of the region bounded by the surface, see [10,17,19].

For a sufficiently dense finite subset of the medial axis transform, the boundary of the union of the corresponding maximal balls, forms a good approximation of the surface. Fig. 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

* Corresponding author. E-mail addresses: nico@cs.rug.nl (N. Kruithof); gert@cs.rug.nl (G. Vegter).

Fig. 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 Fig. 1a.

enclosed by W and the approximating skin surface tends to zero.

In Section 2 we review the theory of skin surfaces as presented in Ref. [13] 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 $O(N^2 \times \log N)$ 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. 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. For the proofs in Sections 4–6 we impose some further constraints on the set of maximal balls. In Section 7 we show that the Power Crust algorithm [3] can be used to construct a sample of maximal balls satisfying these constraints.

1.1. Related work

Amenta et al. [3] show that the medial axis transform of a surface can be effectively approximated by a sufficiently dense sample of points on the surface. If the sample is sufficiently dense, the boundary of the set of approximating

maximal balls is homeomorphic to the original surface. In Ref. [3] Amenta et al. use the approximation of the medial axis transform to compute the power crust, which approximates the original surface by piecewise linear patches.

Another method for visualizing molecules uses Molecular surfaces; See, e.g. [11]. Blobby objects or metaballs [4] also control a surface based on a set of balls. A disadvantage of these surfaces over skin surfaces is that for metaballs there is no known combinatorial structure that decomposes the surface into manageable pieces. Moreover, it is not evident whether and how the input of a metaball surface relates to the medial axis of the surface.

For further reading about skin surfaces we refer to [13] and to [8,9] for background on morphing skin surfaces. In Ref. [7] 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 algorithm can not be used to triangulate our approximating surfaces.

Boissonnat and Cazals [5] 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

This section is a summary of Ref. [13]. For further reading on skin surfaces we refer to this paper.

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 shrunk, and patches of spheres and hyperboloids 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 figures show skin curves in the plane.

Fig. 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.

2.1. The space of weighted points

A 3D weighted point $\hat{p}$ is a pair $\hat{p} = (p, P)$ with center $p \in \mathbb{R}^3$ and weight $p \in \mathbb{R}$. A ball with center p and radius R corresponds to a weighted point (p, R^2) . This definition is slightly different from Ref. [13] 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 pseudo-distance function:

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

for $\hat{p} = (p, P)$ and $\hat{q} = (q, Q)$. 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 Eq. (1) by taking $q = x$ and $Q = 0$. All unweighted points with nonpositive distance to $\hat{p}$ form the ball with center p and radius $\sqrt{P}$. We will make no distinction between the weighted point and this ball. Two weighted points with zero distance are *orthogonal* by definition. Geometrically this means that the spheres bounding two weighted points intersect in a right angle if the two weighted points are orthogonal. An *orthosphere* of a set of weighted points $\mathcal{P}$ is, by definition, a sphere orthogonal to each of the weighted points in $\mathcal{P}$.

2.1.1. General position

In the remainder of this paper we assume general position, by which we mean that no 5 weighted points are equidistant to a point in $\mathbb{R}^3$ and no $k + 2$ centers of weighted points lie on a common k -flat for $k = 0, 1, 2$. We also assume that no four spheres bounding balls in the input set touch in a single point. Several methods like [14] exist to symbolically perturb a data set and ensure these conditions. Note that, under this generality condition, an orthosphere only exists if $|\mathcal{P}| \leq 4$.

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

$$\Pi(\hat{p}) = (\xi_1, \xi_2, \xi_3, \|p\|^2 - P), \quad (2)$$

with $p = (\xi_1, \xi_2, \xi_3)$. Using this vector space we define the addition of two weighted points and the multiplication of a weighted point by a scalar as

$$\hat{p} + \hat{q} = \Pi^{-1}(\Pi(\hat{p}) + \Pi(\hat{q})), \quad c \cdot \hat{p} = \Pi^{-1}(c \cdot \Pi(\hat{p})).$$

Combining these two operators, we introduce the notion of *shrinking* and *growing* of weighted points. Given a shrink factor $s \in \mathbb{R}$, we associate with a weighted point $\hat{p}$ the shrunk weighted point $\hat{p}^s$ defined as follows. Let $\hat{p}' = (p, 0)$. Then

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

A simple calculation shows that $\hat{p}^s = (p, s \cdot P)$

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 .

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

The body of a skin surface $\text{bdy}^s \mathcal{P}$ and skin surface $\text{skn}^s \mathcal{P}$, associated with a set $\mathcal{P}$ of weighted points, are defined by

$$\text{bdy}^s \mathcal{P} = \cup(\text{conv} \mathcal{P})^s \quad (3)$$

$$\text{skn}^s \mathcal{P} = \partial \text{bdy}^s \mathcal{P}. \quad (4)$$

Here $\text{conv} \mathcal{P} \subset \mathbb{R}^3 \times \mathbb{R}$ 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 union of the set of balls. For a skin curve in 2D associated with two weighted points see Fig. 3.

Two sets are homeomorphic (or have the same topology) if there exists a continuous map with a continuous inverse between the sets, see Ref. [15, Ch. 28]. One important feature of a skin surface is that varying the shrink factor does not change the topology, i.e. for $0 < s, s' \leq 1$, $\text{skn}^s \mathcal{P}$ and $\text{skn}^{s'} \mathcal{P}$ are homeomorphic. 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.

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. Later in this section we will show how this is done.

To determine these patches, we need to introduce the mixed complex, which subdivides the skin surface into patches of degree two. The mixed complex is an intermediate structure between the weighted Delaunay triangulation (or: regular triangulation) and its dual, the weighted Voronoi diagram (or: the power diagram). For further reading on Voronoi diagrams we refer to Ref. [6, Ch. 18].

2.2. 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—with respect to the pseudo-distance π defined in Eq. (1)—to $\hat{p}$ than to any other point in $\mathcal{P}$. Formally:

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

A weighted Voronoi cell is a convex polyhedral region, possibly unbounded, or empty. Furthermore, a weighted Voronoi cell $V_{\hat{p}}$ associated with the weighted point $\hat{p} = (p, P)$ does not necessarily contain p .

With a subset $\mathcal{X} \subseteq \mathcal{P}$ we associate the region $\nu_{\mathcal{X}}$ defined by

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

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 $|\mathcal{X}| \leq 4$. In the non-degenerate case, every 0-cell is a point, which is the intersection of four adjacent Voronoi cells. A 1-cell is a line segment, possibly unbounded. And a facet (2-cell) is the intersection of 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 $\hat{p} \in P$. Note that generically $\nu_{\mathcal{X}}$ is determined by $|\mathcal{X}| - 1$ independent equalities (and a number of inequalities). Therefore, $\dim \nu_{\mathcal{X}} = 3 - (|\mathcal{X}| - 1) = 4 - |\mathcal{X}|$.

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

$$\text{Vor}\mathcal{P} = \{\nu_{\mathcal{X}} \mid \mathcal{X} \subseteq \mathcal{P} \wedge \nu_{\mathcal{X}} \neq \emptyset\}.$$

2.3. 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 simplices. A *weighted Delaunay $(3 - \ell)$ -cell* exists for every non-empty weighted Voronoi ℓ -cell $\nu_{\mathcal{X}} \in \text{Vor}\mathcal{P}$ and is the convex hull of the centers of the points in $\mathcal{X}$. Denoting this cell by $\delta - \mathcal{X}$, we have $\delta_{\mathcal{X}} = \{p \mid \hat{p} \in \mathcal{X}\}$. The *weighted Delaunay triangulation* is the subdivision generated by these cells:

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

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 a single point, the *focus* f of $\mathcal{X}$, see [18]. Since the focus lies on the affine hull of the weighted Voronoi cell, it follows that $\pi(\hat{p}, f) = \pi(\hat{q}, f) \forall \hat{p}, \hat{q} \in \mathcal{X}$. The focus corresponding to Delaunay edges are visualised in Fig. 4.

2.4. The mixed complex

The *mixed complex*, associated with a scalar $s \in [0, 1]$, is obtained by taking Minkowski sums of shrunk cells of

Fig. 4. The Delaunay triangulation (solid) and the Voronoi diagram (dashed) of a set of four weighted points (circles). The foci of the edges are visualised by 'x'-symbols.

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

$$\mu_{\mathcal{X}}^s = s \cdot \nu_{\mathcal{X}} \oplus (1 - s) \cdot \delta_{\mathcal{X}},$$

where $\cdot$ denotes the multiplication of a set by a scalar and $\oplus$ the Minkowski sum. 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 polyhedra, it follows that a mixed ℓ -cell is also a convex polyhedron.

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

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

$$\text{Mix}^s\mathcal{P} = \{\mu_{\mathcal{X}}^s \mid \nu_{\mathcal{X}} \in \text{Vor}\mathcal{P}\}.$$

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

2.5. Quadratic patches of a skin surface

The body of a skin surface is by definition the union of the shrunk convex hull of its weighted points $\mathcal{P}$, see Eq. (3). We can extend this definition in the following way:

$$\text{bdy}^s\mathcal{P} = \bigcup \{(\text{conv}\mathcal{X})^s \mid \mathcal{X} \subseteq \mathcal{P}\}$$

In Ref. [13] 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 properly contained inside these subsets. As a result it follows that each part of the skin surface is generated by the shrunk convex hull of at most 4 points. Furthermore, he shows that a subset $\mathcal{X}$ with $\nu_{\mathcal{X}} \neq \emptyset$

Fig. 5. The skin curve of four weighted points (the dotted circles). Each mixed cell contains parts of an hyperbola or a circle.

generates exactly the part of the skin surface that lies inside the mixed cell $\mu_{\mathcal{X}}^s$.

A two-dimensional example is given in Fig. 5. The dotted circles denote the weighted points in the input set. All rectangles (mixed cells corresponding to the Delaunay edges) contain hyperbolic patches. These hyperbolic patches bound the union of the shrunk 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 shrunk with a factor s . On the other hand, for a mixed 2-cell, the skin curve lies outside the weighted point shrunk with a factor $1 - s$, that is orthogonal to every weighted point corresponding to a vertex of the Delaunay triangle. This is the shrunk 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 we have a coordinate system x_1, x_2, x_3 such 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 $\delta_{\mathcal{X}} \cap skn^s \mathcal{P} = \delta_{\mathcal{X}} \cap S$, 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 pseudo-distance between the focus and a weighted point in the Delaunay cell.

For $k = 0$ and $k = 3$, the mixed cell contains a sphere, where the two cases are distinguished by the fact that, for $k = 0$, the interior of the skin surface lies inside the sphere, whereas for $k = 3$, the exterior of the skin surface lies inside the sphere. For $k = 0$, the mixed cell corresponds to a weighted Delaunay vertex. These mixed cells contain a shrunk weighted point from the input set. For $k = 1$ and $k = 2$ the mixed cell contains a hyperboloid.

For $k = 1$ the rotation axis is aligned along the Delaunay edge and for $k = 2$ it is aligned along the Voronoi edge. These patches of hyperboloids ensure that the transition between the spherical patches is tangent continuous.

For the proof that these patches form the corresponding skin surface we refer to Ref. [13, Theorem 13].

3. Approximation algorithm

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

A maximal ball is a ball contained in a surface W that is maximal with respect to inclusion. We say that a ball $\hat{p}$ is W -maximal if it is a maximal ball of W . The Hausdorff distance between two subsets X and Y of a Euclidean space is $d(X, Y)$ the minimal distance of a point in X to a point in Y . For a formal definition of the Hausdorff distance see Ref. [15, Chapter 43].

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

- (i) the boundary ∂M of M is homeomorphic to W ;
- (ii) every ball in $\mathcal{P}$ is M -maximal;
- (iii) The Hausdorff distance between M and the body enclosed by W is at most ε .

For a given C^2 -surface W and a positive constant ε , the approximation algorithm takes an ε' -admissible set $\mathcal{P}$ of balls, $0 < \varepsilon' < \varepsilon$, and computes a C^1 skin surface S associated with $\mathcal{P}$ satisfying the following properties:

- $C_1(S)$: W and S are homeomorphic;
- $C_2(S)$: Every ball in $\mathcal{P}$ is S -maximal;
- $C_3(S)$: The Hausdorff distance between the bodies of W and S does not exceed ε .

The conditions for an ε -admissible set are rather strong, but not all conditions are needed for the properties to hold. In fact, we only need condition (i) for $C_1(S)$, (ii) for $C_2(S)$ and (iii) for $C_3(S)$. In Section 7 we show that the *power crust algorithm* [3] can be used to generate an ε -admissible set.

From the definition of shrinking, Eq. (3), it follows that $(\hat{p}^{1/s})^s = \hat{p}$. Therefore, $skn^s(\mathcal{P}^{1/s})$ contains the balls in $\mathcal{P}$. We define our approximating skin surface $S_s(\mathcal{P})$ as follows.

Definition 2. 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 $\mathcal{P}^{1/s}$.

If no confusion is possible, we write S_s instead of $S_s(\mathcal{P})$. If $\mathcal{P}$ is an ε -admissible set of balls then all three properties $C_1(S_s)$ – $C_3(S_s)$ are satisfied. Indeed, $\mathcal{P}^1 = \mathcal{P}$, so S_1 is

Fig. 6. The set of maximal circles and the approximating skin curve for different shrink factors (decreasing from left to right).

the boundary of the union of these balls and the properties are trivially satisfied. Obviously, S_1 is *not smooth*. However, for values of s slightly smaller than 1 the properties are satisfied, and S_s is a C^1 -surface. Indeed, for s sufficiently near 1, the surface S_s is homeomorphic to S_1 , and hence to W . Furthermore in general each ball touches their union in a spherical patch. For s near 1, the patches between the input balls are small and therefore all input balls contribute to the skin surface in a spherical patch. These patches make the balls maximal. Since these patches are small, the Hausdorff distance of these patches to $\cup \mathcal{P}$ is also small.

One or more of these conditions may be violated for smaller values of s . To illustrate these violations in the 2D-case we depict a sequence of skin curves corresponding to decreasing values of s in Fig. 6. Fig. 6a 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 Fig. 6b.

Decreasing the shrink factor, i.e. growing these patches, causes a change in the topology of the skin curve, i.e. a violation of condition $C_1(S_s)$, see Fig. 6d. Finally, as we decrease the shrink factor even further, the balls are no longer maximal, i.e. a violation of $C_2(S_s)$, as depicted in Fig. 6e. In Section 4, we show how to compute the time when the change of topology occurs. In Section 5 we derive, how far the shrink factor can be decreased without loss of maximality of the balls.

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

Theorem 3. A value $s_i(\mathcal{P}) \leq 1$ such that for s , with $s_i(\mathcal{P}) < s < 1$, condition $C_i(S_s)$ is satisfied for $i = 1, 2, 3$, can be computed in $O(n^2 \log n)$ time and $O(n^2)$ space, where n is the number of balls in $\mathcal{P}$.

In fact $s_1(\mathcal{P})$ is the minimal value for which the theorem holds. The algorithm presented in Section 4 computes $s_1(\mathcal{P})$ 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 derive some results for growing balls, which are of independent interest. We introduce a growth factor t , and later on we apply the results of this section for $t = 1/s$, where s is the shrink factor.

We compute the smallest t , such that the boundary of the union of the balls in $\mathcal{P}^t$ is not homeomorphic to $\partial(\cup \mathcal{P})$, viz. the transition of Fig. 6c–d. Since the skin surface $S_s(\mathcal{P})$ is homeomorphic to the boundary of the union of the set of balls $\mathcal{P}^t$, see Section 2, this is equivalent to computing the value $s_1(\mathcal{P})$, defined in Section 3. Fig. 7 illustrates this equivalence in 2D, where the change of topology of the skin curve coincides with the change of topology of the union of balls defining the skin curve.

Let $\mathcal{Q}$ be a set of balls in $\mathbb{R}^3$.

Definition 4. The boundary $\partial(\cup \mathcal{Q}^t)$ of the union of balls in $\mathcal{Q}^t$ 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.

Fig. 7. Change of topology at (x, τ) of a union of two (top row) and three (bottom row) weighted points.

The following result states that a set of weighted points changes topology if their intersection becomes non-empty.

Lemma 5. 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. If $\cap \mathcal{Q}^t$ is empty, then it is not for all points to contribute in a change of topology. 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 $\mathcal{Q}$ are increasing with t , whereas the centers are independent of t , it follows that, for $t > t'$, $\cup \mathcal{Q}^t$ is star shaped with respect to y . In particular, for $t > t'$, the set $\partial(\cup \mathcal{Q}^t)$ is non-empty, and does not change topology (in fact $\partial(\cup \mathcal{Q}^t)$ is homeomorphic to a 2-sphere). 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 $t \geq \tau$. $\square$

It follows from the proof that, for $t > \tau$, the union $\cup \mathcal{Q}^t$ is homeomorphic to a disk. Furthermore, x is an interior point of the intersection $\cap \mathcal{Q}^t$.

Lemma 6. If $\partial(\cup \mathcal{P}^t)$ 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 5 that x belongs to the boundary of the weighted points in $\mathcal{Q}^\tau$, 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 $\mathcal{P}^\tau$. $\square$

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 co-dimension $|\mathcal{Q}|$. 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 co-dimension of this intersection is less than $|\mathcal{Q}|$. We impose the following generic condition on the family $\mathcal{P}$.

Generic change of topology. If $\partial(\cup \mathcal{P}^t)$ 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 $\mathcal{Q}^\tau$ intersect in an affine space of co-dimension $|\mathcal{Q}| - 1$.

Fig. 8. The topological changes in a skin surface when the input balls are grown. We omitted the filling of a void, since this happens inside the skin surface. The analogue in 2D of the latter topological change is depicted in the bottom row of Fig. 7.

In the plane, at most three weighted points are involved in a generic change of topology, see Fig. 7. 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 Fig. 8.

To describe the situation in three-space, let $\mathcal{Q}$ be the set of balls involved in a generic change of topology, 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 $|\mathcal{Q}| = 4$.

4.1. Lifting the weighted Voronoi diagram

In view of Lemma 6 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 the extended space $\mathbb{R}^3 \times \mathbb{R}$. 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

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

The boundary of this half space is called the *extended bisector* of the weighted points p and 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 $\hat{p} = (p, P) \in \mathcal{P}$, 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 of scaled weighted points. We assume that the weighted points in $\mathcal{P}$ are in general position in the sense that

1. The centers of $k + 2$ weighted points in $\mathcal{P}$ lie on a common k -flat, for $0 \leq k \leq 2$;
2. no 6-tuple of weighted points in $\mathcal{P}^t$ has a common orthosphere;
3. a 5-tuple of weighted points in $\mathcal{P}^t$ has a common orthosphere for an unique value of t .

For a set of weighted points in general position, there is a constant number of 1-faces and 2-faces incident upon a vertex of the extended Voronoi diagram.

If $\mathcal{P}$ is in general position, then, for $0 \leq k \leq 4$, the k -flat supporting 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 of the extended Voronoi diagram is defined by 5 weighted points, it is incident upon

$$\binom{5}{4} = 5$$

edges. The number of edges incident upon a vertex (x, τ) of the extended Voronoi Diagram, and contained in the half space below (above) the hyperplane $t = \tau$, is at least 1 since the extended Voronoi diagram subdivides $\mathbb{R}^3 \times \mathbb{R}$. If a vertex has $k + 1$ incident edges in the half space below the hyperplane $t = \tau$, $0 \leq k \leq 3$, then this vertex is said to be 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 7. *The extended Voronoi Diagram of a set of N weighted points in general position 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 we assume general position, the total number of faces of the extended Voronoi Diagram is $O(N^2)$. $\square$

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

Fig. 9. A set of weighted points for which the Voronoi diagram changes topology a quadratic number of times. (a) shows the set of weighted points and (b) its Voronoi diagram. (c) and (d) show the Voronoi diagram for increasing t .

We construct this set on a plane. First align $N/2$ points along the vertical axis with a constant weight P . The other weighted points are aligned along the horizontal axis and have a weight P' , with $P' < P$. Ensure that the Voronoi cells of two of these weighted points touch the Voronoi cells of all other weighted points along the vertical axis by moving them away from the origin. 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 Fig. 9a, and its Voronoi diagram in Fig. 9b.

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 horizontal 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 become larger, whereas the Voronoi cells of the weighted points on the horizontal axis become smaller. 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 Fig. 9b–d.

Proposition 8. *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 with planes $t = \text{constant}$, starting at $t = 0$. The weighted Voronoi Diagram for $t = 0$ can be computed in $O(N^2)$ time, using the algorithm from [16]. 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 by definition 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 parameterized 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;
- (ii) Update the priority queue by adjusting the destruction times of all j -simplices, $j > 0$, incident with the newly created simplex.

Adding the destruction time of the newly created simplex is not necessary since it has the value of the current event. 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 7 the total number of time stamps is $O(N^2)$. Therefore the time complexity of the algorithm is $O(N^2 \log N)$. $\square$

In practice, there is no need to compute the entire extended Voronoi diagram and our algorithm terminates as soon as one of the properties of our approximation is violated.

Corollary 9. Let $\mathcal{P}$ be a set of N weighted points in $\mathbb{R}^3$. Then τ_0 , defined by

$$\tau_0 = \min\{\tau > 1 \mid \partial(\cup \mathcal{P}^t) \text{ changes topology for } t = \tau\},$$

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 a unique time stamp $\tau - \mathcal{Q}$ is defined for 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. $\square$

5. Preserving maximality of the balls

The medial axis transform is a representation of a surface (the boundary of the union of all maximal balls is

the surface). An ε -admissible set approximates the medial axis transform of a surface W by a set $\mathcal{P}$ of approximate maximal balls of W . In this section, we compute a shrink factor $s_2(\mathcal{P})$, such that all balls in $\mathcal{P}$ are maximal in the approximating skin surface $S_s(\mathcal{P})$, for $s_2(\mathcal{P}) < s \leq 1$. This means that we do not discard any information given by the maximal balls in $\mathcal{P}$.

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

Requirement (ii) of an ε -admissible set ensures that each weighted point generates at least one spherical patch on ∂M , otherwise the weighted point is not ∂M -maximal. In Fig. 10, we show that for any $r > 0$, we can construct an ε -admissible set that contains a weighted point $\hat{p}$ such that $\hat{q}$ generates only one circular arc and this arc is arbitrarily 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 such a 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 separately. We stop decreasing s just before some spherical patch touches S_s only in one point.

To determine when a spherical patch 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

Fig. 10. Three weighted points in an ε -admissible set $\mathcal{P}$. The radius of the maximal circle $\hat{p}$ in the middle is shrunk with a factor ε . The result is that $\hat{p}$ only touches $\cap \mathcal{P}$ in one small circular arc.

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 $\hat{p}$ shrunk with a factor s towards the weighted Delaunay 0-cell $\delta_{\{\hat{p}\}}$.

Lemma 10. *Let γ be a spherical patch on the boundary of $\cup \mathcal{P}$ generated by a weighted point $\hat{p} \in \mathcal{P}$. Then there is a weighted Delaunay 3-cell δ_x adjacent to $\delta_{\{\hat{p}\}}$ for which the segment between p and v_x intersects γ . Furthermore, the skin surface S_s contains a spherical patch of γ iff*

$$s \cdot |p - v_x| > P$$

Proof. First, we analyze 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 shrunk 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 of the adjacent Voronoi 0-cells.

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

$$s \cdot v_x + (1 - s) \cdot \delta_{\{\hat{p}\}}.$$

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 v_x also depends on the shrink factor. $\square$

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. Through our study of degeneration, 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 prove the following approximation result.

Proposition 11. *Let $\varepsilon > 0$, and let $\mathcal{P}$ be an ε' -admissible sample of a surface W , where $0 < \varepsilon' < \varepsilon$. Then there is an*

$s_3(\varepsilon, \varepsilon')$, with $0 < s_3(\varepsilon, \varepsilon') < 1$, such that

1. *$s_3(\varepsilon, \varepsilon')$ can be computed in $O(N^2 \log N)$ time;*
2. *For $s_3(\varepsilon, \varepsilon') \leq s \leq 1$, the Hausdorff-distance between the body B_W bounded by W and the body $B_S(s)$ bounded by $S_s(\mathcal{P})$ does not exceed ε .*

Since $\mathcal{P}$ is ε' -admissible it follows that $d(B_W, \cup \mathcal{P}) \leq \varepsilon'$, see Section 3, condition (iii). Therefore, the main result would be clear if we can prove that $d(\cup \mathcal{P}, B_S(s)) \leq \varepsilon - \varepsilon'$ for shrink factors $s_3(\varepsilon, \varepsilon') \leq s < 1$. The latter statement follows from the following result by taking $s_3(\varepsilon, \varepsilon') = s_3'(\varepsilon - \varepsilon')$:

Lemma 12. *For $\eta > 0$ there is an $s_3'(\eta)$, with $0 < s_3'(\eta) < 1$, such that*

1. *$s_3'(\eta)$ can be computed in $O(N^2 \log N)$ time;*
2. *For $s_3'(\eta) \leq s \leq 1$, the Hausdorff-distance between $\cup \mathcal{P}$ and $B_S(s)$ does not exceed η .*

To prove the lemma, we need three technical lemmas. The first result states that, for an efficiently computable range of shrink factors, the body $\mathcal{B}_\mathcal{P}(s) = \text{bdy}^s(\mathcal{P}^{1/s})$ of the skin surface $S_s(\mathcal{P})$ intersects a cell μ_x^x of the mixed complex of $\mathcal{P}^{1/s}$ only if the focus $f(\mathcal{X})$ of $\mathcal{X}$ is contained in $\cup \mathcal{X}$.

Lemma 13. *There is an s_3'' , with $0 < s_3'' < 1$, such that*

1. *s_3'' can be computed in $O(N^2 \log N)$ time;*
2. *If $s_3'' \leq s \leq 1$ and $\mu_x^s \cap \mathcal{B}_\mathcal{X}(s) \neq \emptyset$, then $f(\mathcal{X}) \in \cup \mathcal{X}$.*

Proof. The main idea of the proof is again to use the sweep-line algorithm and to compute the maximal shrink factor for which $\mu^s - \mathcal{X} \cap \mathcal{B}_\mathcal{X}(s) \neq \emptyset$ locally. By taking the maximum of all these local values, we obtain a global maximum. This global maximum is the minimal value for which condition (2) is not satisfied.

In the remainder of this proof we show how to compute the shrink factor locally for different types of mixed cells. The focus of a mixed 0-cell is the center of the generating weighted point, which lies inside the weighted point.

A mixed 1-cell corresponding to a set $\mathcal{X}$ contains 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 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 $f(\mathcal{X})$ and the center of a weighted point in $\mathcal{X}$. Moreover, the skin surface

intersects the mixed cell if

$$s \cdot |p - f(\mathcal{X}^{1/s})| \leq P, \quad (5)$$

for some $\hat{p} \in \mathcal{X}$. This can be shown in a similar way as done in the proof of Lemma 10 and is left as an exercise to the reader. $\square$

As a corollary, for shrink factors in the range $[s_3'', 1]$ the patches of the skin surface are defined by subsets $\mathcal{X}$ of $\mathcal{P}$ for which all weighted points in $\text{aff}(\mathcal{X})$ have positive weights. This is made more precise in our second technical lemma.

Lemma 14. *If $f(\mathcal{X}) \in \cup \mathcal{X}$, then the weight P of a weighted point $\hat{p} \in \mathcal{X}$ is non-negative.*

Proof. A simple, but tedious, calculation shows that for $\hat{p} \in \text{aff}(\mathcal{X})$, $\pi(\hat{p}, f(\mathcal{X})) = \text{constant}$. Since $f(\mathcal{X})$ lies inside $\hat{p}' \in \mathcal{X}$, $\pi(\hat{p}', f(\mathcal{X})) \leq 0$, which means that $\pi(\hat{p}, f(\mathcal{X})) \leq 0$. As a result, $\|p - f(\mathcal{X})\|^2 \leq P$, so $P \geq 0$. Therefore, all weighted points in $\text{aff}(\mathcal{X})$ have positive weight. $\square$

Our third technical result provides us with an upper bound on the Hausdorff-distance between a quadratic patch and (a subset of) the union of balls in the input set. Let $\mathcal{X}$ be a subset of $\mathcal{P}$ with $f(\mathcal{X}) \in \cup \mathcal{X}$. Let $X = \{p | \hat{p} \in \mathcal{X}\}$ and let $m \in \text{aff}(X)$ be the point equidistant to each $x \in X$. Let d_x be the corresponding distance $\|m - x\|$.

Lemma 15. *For $s_{\mathcal{X}}(\varepsilon) \leq s \leq 1$ with*

$$s_{\mathcal{X}}(\varepsilon) = 1 - \frac{\varepsilon^2}{d_x^2}, \quad (6)$$

we have that

$$d(\cup \mathcal{X}, \mathcal{B}_{\mathcal{X}}(s)) \leq \varepsilon.$$

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

The set $\Pi(\text{conv} \mathcal{X})$ forms a hyperplane in $\mathbb{R}^4$. Since the weight is the distance along the last coordinate axis to the unit paraboloid, $P_1(p)$ 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 Fig. 11. The top parabola denotes the set of weighted points with zero weight.

A similar reasoning shows that $\Pi(\text{conv} \mathcal{X}^{1/s})$ is also a hyperplane in $\mathbb{R}^4$. By the definition of shrinking, $P_s(p)$ is a paraboloid with leading coefficient s .

Knowing the leading coefficient of $P_s(p)$, we now determine its top. Consider a weighted point $\hat{p}' \in \mathcal{X}$. It is first shrunk with a factor $1/s$ and then with a factor s , therefore $P_s(p') = P'$, as is also depicted in Fig. 11. Using these points the parabola is uniquely defined, and it

Fig. 11. The vector space obtained under the lifting operator Π , defined in Eq. (2), 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.

follows that

$$P_s(p) = P_1(p) - (1 - s)(\|x - m\|^2 - \|p - m\|^2),$$

for some $x \in X$, since $\|x - m\| = \|x' - m\|$ for $x, x' \in X$. Indeed, $P_s(p)$ is a paraboloid with leading coefficient s , since $P_1(p)$ has leading coefficient 1. Further, the weight of the points in $\mathcal{X}$ is independent of the shrink factor.

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

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

The equation in the lemma follows if we take $p = m$ to obtain the maximal difference in weight, and then solve $d(\hat{p}(s), \cup \mathcal{P}) = \varepsilon$ 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 $p \in \text{conv} X$ $\square$

Proof of Lemma 12. Since $\cup \mathcal{P} \subset B_S(s)$, for $0 < s \leq 1$, it follows that $d(x, B_S(s)) = 0$ for $x \in \cup \mathcal{P}$. Conversely, let $x \in B_S(s)$. Let $s_3'(\eta)$ be the maximum of s_3' and

$$\max\{s_{\mathcal{X}}(\eta) | f(\mathcal{X}) \in \cup \mathcal{X}\},$$

then the statement follows from Lemmas 13–15.

7. Generating input data

This section shows one way to obtain an ε -admissible set of maximal balls from a finite sample of points on the surface W bounding a volume B_W . In the computational geometry community several algorithms [1,3,12] have been proposed to reconstruct a surface from a finite set of sample points on the surface. Some of these algorithms first approximate the medial axis (transform) and use it for reconstruction. In this section, we show that the approximate medial axis transform of the power crust [3]

can also be used to generate an ε -admissible sample of maximal balls.

The extended medial axis transform is the medial axis transform not restricted to the inside of the surface. All maximal balls in the extended medial axis transform touch the surface in at least two points, and contain no surface points in their interior. The extended medial axis is the set of centers of the maximal balls in the extended medial axis transform.

7.1. The power crust

We will only briefly discuss the power crust here and refer to Ref. [3] for a more detailed discussion. The method is based on the observation that the Voronoi cells of a dense sample X on a surface W are elongated in the direction of the normals.

To be more precise, the *poles* of a sample point x are the two points in the Voronoi cell of x farthest away from x , such that the two points lie on different sides of W . The balls centered at the poles that touch the sample are called the *ph-polar balls*. The main observation of Amenta and Bern [1] is that the two poles lie close to the extended medial axis, one to the interior and the other to the exterior medial axis.

When the polar balls are found, it remains to label the balls in- and outside. Therefore, consider the Voronoi diagram of the polar balls. Start with labeling one polar ball on the convex hull as outside and label adjacent polar balls according to the following three criteria: first, two polar balls corresponding to the same sample have opposite labels. Secondly, if the two balls have a shallow intersection, they lie on opposite sides of the surface, because of the assumption that distance between sample points is small with respect to the radius of the polar balls. On the other hand, if two balls intersect deeply, they have the same label.

To complete the surface reconstruction process, the Voronoi facets separating inner from outer poles are returned.

7.2. Sampling of the surface

Under some conditions on the sample of points on the surface, the power crust can be proven to give accurate results. The following definitions are taken from [1,2].

Definition 16. *The Local feature size (LFS) at a point $w \in W$ is the distance from w to the nearest point of the extended medial axis.*

The LFS is a measure of how close another part of the surface is to a surface point. Note that the LFS of x can be smaller than the minimal distance to the centers of the maximal balls touching in x .

Definition 17. *For $r > 0$, a set $X \subseteq W$ is an r -sample if the distance from any point $w \in W$ to its closest neighbor in X is at most $r \cdot \text{LFS}(w)$*

For an r -sample with $r \leq 0.1$, the algorithm labels the polar balls correctly. Furthermore, the union of the inner polar balls is homeomorphic to B_W . For the proofs in the next paragraph we assume that this sampling condition is satisfied.

7.3. Construction of an ε -admissible set

We will use the power crust to obtain an ε -admissible set from an r -sample X of sample points on a surface W . Let P be the set of inner polar balls and $M = \cup P$.

Condition (i) (see page 5) is proven directly in Ref. [3 Lemma 25], that states that there exists a homeomorphism between ∂M and W . For condition (ii), we prove the following lemma.

Lemma 18. *All inner polar balls $\mathcal{P}$ are maximal balls in M .*

Proof. The maximal balls in $\mathcal{P}$ are polar balls of X . These polar balls are centered at Voronoi vertices of X . Therefore, each $\hat{p} \in \mathcal{P}$ touches four points in X . Since all points in X lie on ∂M [3, Observation 3] all input balls are maximal. $\square$

Although all balls in $\mathcal{P}$ are maximal in ∂M , there can be a set of balls $\mathcal{P}' \subset \mathcal{P}$ that touch ∂M only in points of X . For any $s < 1$, smooth patches grow over the points in X and the balls in $\mathcal{P}'$ are not maximal anymore. Since the balls in $\mathcal{P}'$ are degenerate maximal balls, we do not consider them when we compute the minimal shrink factor for which all balls are maximal.

We prove the (iii) using the tubular neighborhood of the surface W , cf. Ref. [3].

Definition 19. *The tubular neighborhood around the surface W is the set of points within distance $r \cdot \text{LFS}(w)$ of a point $w \in W$*

We can now prove criterion (iii).

Lemma 20. *The Hausdorff distance $d(B_W, M)$ is $O(r)$.*

Proof. The bodies M and B_W only differ in the tubular neighborhood T of W , since all poles are labeled correctly ([3, Lemma 37]) and their boundaries lie in the tubular neighborhood ([3, Theorem 21] and the definition of T). Consider a point $m \in M \cap T$, the distance of m to B_W/T is $O(r)(w)$ for some $w \in W$. We bound $\text{LFS}(w)$ by $\max_{w' \in W} \times (w')$, which is a constant since W does not change. Thus

the distance of m to B_W is $O(r)$. A similar reasoning holds for the distance between a point $w \in B_W \cap T$ and M . $\square$

To conclude, the power crust constructs an ε -admissible set of maximal balls from an r -sample on a surface, where $\varepsilon = O(r)$. For small ε this means that there is a constant c such that $\varepsilon \leq c \cdot r$. Therefore, the power crust can construct an ε -admissible set from an ε/c -sample. The constant c can be rather large. This is due to the fact that the Hausdorff distance is a global measure between the two surfaces, whereas the LFS is defined locally on the surface.

8. Conclusion and future work

We presented an algorithm to effectively compute a C^1 -approximation S of a C^2 -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 the interpolating patches between balls are always concave. Another drawback of our algorithm is that it determines the shrink factor globally: if a high shrink factor is needed to satisfy some condition in a small part of the skin surface, this influences the approximation of the whole surface. Maybe it is possible to use different values of the shrink factor in different parts of the surface, according to some criterion.

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 hardly improves on the union of the balls. We assume that the shrink factor will be significantly smaller than 1. Let s_W be the shrink factor obtained by our algorithm for the medial axis transform of W , we conjecture:

Conjecture 21. *For a C^2 -surface W there is an $s_W < 1$ such that the shrink factor s_S corresponding to a sample S 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 LFS 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.

Acknowledgements

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).

References

- [1] Amenta N, Bern M. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* 1999;22(4):481–504.
- [2] Amenta N, Bern M, Eppstein D. The crust and the β -skeleton: Combinatorial curve reconstruction. *Graph Models Image Process* 1998;60:125–35.
- [3] Amenta N, Choi S, Kolluri RK. The power crust, unions of balls, and the medial axis transform. *Comput Geom Theory Appl* 2001;19: 127–53.
- [4] Blinn J. A Generalization of Algebraic Surface Drawing. *ACM Trans Graph* 1982;1(3):235–56.
- [5] Boissonnat J-D, Cazals F. Smooth surface reconstruction via natural neighbour interpolation of distance functions. *Comp Geometry Theory Appl* 2003; in press.
- [6] Boissonnat J-D, Yvinec M. *Algorithmic geometry*. UK: Cambridge University Press; 1998. Translated by Hervé Brönnimann.
- [7] Cheng H-L, Dey TK, Edelsbrunner H, Sullivan J. Dynamic skin triangulation. *Discrete Comput Geom* 2001;25:525–68.
- [8] Cheng H-L, Edelsbrunner H, Fu P. Shape space from deformation. *Comput Geom Theory Appl* 2001;19:191–204.
- [9] Cheng S-W, Edelsbrunner H, Fu P, Lam K-P. Design and analysis of planar shape deformation. *Comput Geom Theory Appl* 2001;19: 205–18.
- [10] Choi HI, Choi SW, Moon HP. Mathematical theory of medial axis transform. *Pacific J Math* 1997;181(1):57–88.
- [11] Connolly ML. Analytical molecular surface calculation. *J Appl Crystallogr* 1983;16(5):548–58.
- [12] Dey TK, Goswami S. Tight cocone: a water-tight surface reconstructor. *Proceedings of the eighth ACM symposium on Solid modeling and applications*, ACM Press; 2003. p. 127–134.
- [13] Edelsbrunner H. Deformable smooth surface design. *Discrete Comput Geom* 1999;21:87–115.
- [14] Edelsbrunner H, Mücke EP. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans Graph* 1990;9(1):66–104.
- [15] Goodman JE, O'Rourke J. In: Goodman JE, O'Rourke J, editors. *Handbook of discrete and computational geometry*. Boca Raton, FL: CRC Press LLC; 1997.
- [16] Joe B. Construction of three-dimensional Delaunay triangulations using local transformations. *Comput Aided Geom Design* 1991;8(2): 123–42.
- [17] Lieutier A. Any open bounded subset of $\mathbb{R}^n$ has the same homotopy type than its medial axis. *Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications*, ACM Press; 2003. pp. 65–75.
- [18] Pedoe D. *Geometry, a comprehensive course*. New York: Dover Publications; 1970.
- [19] F. Wolter, Cut locus and medial axis in global shape interrogation and representation. Technical report, December 1993.

Gert Vegter did his PhD in Mathematics (1983) at the University of Groningen. Currently he is an adjoint professor at the Department of Mathematics and Computing Science at the University of Groningen. His research interests include computational geometry and topology, computer aided geometric design, and algorithmic aspects of dynamical systems.

Nico Kruithof received a Masters degree in Computer science from the University of Groningen and is currently a PhD student at the University of Groningen. His research interests include computational geometry and topology and computer aided geometric design, especially in applications of Delaunay triangulations and Skin Surfaces.