Meshing Skin Surfaces with Certified Topology

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

Abstract

Skin surfaces are used for the visualization of molecules. They form a class of tangent continuous surfaces defined in terms of a set of balls (the atoms of the molecule) and a shrink factor. More recently, skin surfaces have been used for approximation purposes.

We present an algorithm that approximates a skin surface with a topologically correct mesh. The complexity of the mesh is linear in the size of the Delaunay triangulation of the balls, which is worst case optimal.

We also adapt two existing refinement algorithms to improve the quality of the mesh and show that the same algorithm can be used for meshing a union of balls.

Keywords: Skin Surfaces, Meshing, Isotopy, Regular triangulation.

1 Introduction

Skin surfaces, introduced by Edelsbrunner in [14], have a rich and simple combinatorial and geometric structure that makes them suitable for modeling large molecules in biological computing. Meshing such surfaces is often required for further processing of their geometry, like in numerical simulation and visualization. We present an algorithm for meshing skin surfaces with guaranteed topology.

Large molecules can be modeled using skin surfaces by representing each atom by a sphere. Atoms that lie close to each other are connected by smooth patches. A skin surface is parameterized by a set of weighted points (input balls) and a shrink factor. If the shrink factor is equal to one, the surface is just the boundary of the union of the input balls. If the shrink factor decreases, the skin surface

*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).

becomes tangent continuous, due to the appearance of patches of spheres and hyperboloids connecting the balls.

We present an algorithm in [20] that approximates an arbitrary surface with a skin surface. The approximation is homeomorphic to the skin surface and the Hausdorff distance between the two surfaces is arbitrarily small.

Two surfaces embedded in three space are isotopic if there is a continuous deformation within the embedding space that transforms one surface into the other one. In particular, isotopic surfaces are homeomorphic. The algorithm presented in this paper constructs a mesh isotopic to the skin surface in two steps: it constructs a coarse, isotopic mesh which is subsequently improved by slightly adapted refinement algorithms. The complexity of the coarse mesh is quadratic in the number of input balls, and is independent of the shrink factor. This is worst case optimal. For the second step a broad range of refinement algorithms can be used. Existing algorithms may have to be adapted slightly to ensure the isotopy. We show how this is done for the refinement algorithms of Chew [11] and Kobbelt [18]. The $\sqrt{3}$-subdivision algorithm by Kobbelt is very fast, and refines the size of the triangles. However, it does not improve the quality of the mesh elements in terms of angle size. Chew's algorithm improves the quality of the mesh in terms of the angles and size of the triangles. The quality mesh is suitable for numerical computations. Our version of these algorithms preserve the isotopy property.

Related work. Most existing algorithms meshing implicit surfaces do not guarantee topological equivalence of the surface and the mesh constructed. The marching cubes algorithm [21] subdivides a region into cubes and triangulates the surface within these cubes based on whether the vertices of the cube lie inside or outside the cube. A variant of this algorithm that follows the surface is [5].

The marching triangulation method [17] extends a small initial mesh by walking over the implicit surface, starting from a seed point. Our paper [19] presents a marching triangulation method for meshing skin surfaces by carefully choosing the step size during the walk over the mesh. However, as the shrink factor goes to one or to zero, the size of the mesh goes to infinity.

The algorithms in [9, 16] construct a topologically correct mesh approximating a skin surface in the special case of a shrink factor 0.5. It is likely that this algorithm can be generalized to work for arbitrary shrink factors, but this would probably result in a denser mesh in order to guarantee the topology. The algorithm is also rather slow. Another approach is found in [10]. We could not verify the claim that the mesh produced by the algorithm in this paper is homeomorphic to the skin surface. The idea is to approximate the Morse-Smale complex of the skin surface and use a marching algorithm to approximate the skin surface using this Morse-Smale complex. There is no guarantee that the approximating

Morse-Smale complex is topologically correct since the computation of the separatrices requires solving an ordinary differential equation using a Runge-Kutta method.

Another method for visualizing molecules uses Molecular Surfaces [13]. Visualization algorithms for this type of surfaces are presented in [3, 4]. The algorithms presented in [6, 23] are the first general methods guaranteeing topological equivalence of the implicit surface and the mesh.

Contribution. The approach to meshing skin surfaces described in this paper is new. The main contribution compared to [9] is that our approach works for any shrink factor. We also establish isotopy, which is stronger than topological equivalence. Our algorithm is more flexible in the sense that we generate a coarse mesh that is isotopic to the skin surface and can be refined by different algorithms, as shown in Section 3.4, whereas the algorithm in [9] immediately constructs a homeomorphic quality mesh. Further, our algorithm is much faster. It constructs a mesh in minutes where the algorithm presented in [9] takes hours.

On the theoretical side, we analyze the structure of the mixed complex and decompose the mixed cells into tetrahedra. Within a tetrahedron the intersection with the skin surface is either empty or a topological disk. It is fairly easy to extract the isotopic mesh from this tetrahedral complex by a marching tetrahedra algorithm.

Outline. In Section 2 we extend the theory of skin surfaces as presented in [14]. We start by introducing a hierarchical conbinatorial structure on the mixed complex. With each face of this complex we associate an anchor point, which plays a crucial role in the meshing algorithm. Section 3 describes the construction of the coarse mesh and establishes the isotopy between this mesh and the skin surface. In Section 3.4, we describe the two methods to improve the coarse mesh (i) subdivision of the triangles and (ii) improvement of the quality of the triangles with regard to the size of the minimal angle. Finally, we describe our implementation and give results in Section 4 and 5.

2 Definitions

This section first briefly reviews skin surfaces and then introduces some new concepts specific to the meshing algorithm. For a more thorough introduction to skin surfaces, we refer to [14] where they were originally introduced.

2.1 Skin surfaces.

A skin surface is defined in terms of a finite set of weighted points $\mathcal{P}$ and a shrink factor s, with $0 \leq s \leq 1$. A weighted point $\hat{p} = (p, \mathrm{P}) \in \mathbb{R}^d \times \mathbb{R}$ corresponds

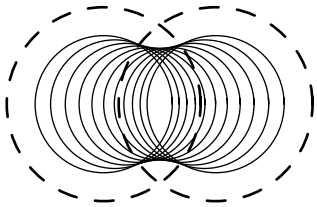

Figure 1: The skin curve of two weighted points (the two dashed circles). The smaller circles form a subset of the shrunken convex hull of the input points. Its boundary forms the skin curve.

to a ball with center p and radius $\sqrt{\mathrm{P}}$. A pseudo-distance between two weighted points is given by:

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

where $\hat{p} = (p, \mathrm{P})$, $\hat{q} = (q, \mathrm{Q})$ and $\| \cdot \|$ denotes the Euclidean distance. The pseudo-distance $\pi(\hat{p}, x)$ of a weighted point $\hat{p}$ to an (unweighted) point x is the pseudo-distance of $\hat{p}$ to the weighted point centered at x with zero weight. Two weighted points with zero distance are called *orthogonal*. An *orthosphere* of a set of weighted points $\mathcal{P}$ is, by definition, a sphere orthogonal to each of the weighted points in $\mathcal{P}$.

The space of weighted points inherits a vector space structure from $\mathbb{R}^{d+1}$ via the bijective map $\Pi : \mathbb{R}^d \times \mathbb{R} \to \mathbb{R}^{d+1}$, defined by $\Pi(\hat{p}) = (x_1, \ldots, x_d, \| p \|^2 - \mathrm{P})$, with $p = (x_1, \ldots, x_d)$. Addition of two weighted points and the multiplication of a weighted point by a scalar are defined in the vector space structure inherited under Π. For further reading on the space of circles and spheres we refer to [22].

Starting from a weighted point $\hat{p} = (p, \mathrm{P})$, the shrunken weighted point $\hat{p}^s$ is defined as $\hat{p}^s = (p, s \cdot \mathrm{P})$. The set $\mathcal{P}^s$ is the set obtained by shrinking every weighted point of $\mathcal{P}$ by a factor s.

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

$$\mathtt{bdy}^s\mathcal{P} = \cup(\mathtt{conv}\ \mathcal{P})^s \tag{2}$$

$$\mathtt{skn}^s\mathcal{P} = \partial\ \mathtt{bdy}^s\mathcal{P}. \tag{3}$$

Here $\mathtt{conv}\ (\mathcal{P}) \subset \mathbb{R}^d \times \mathbb{R}$ 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 $\mathbb{R}^d$ – of the union of the corresponding set of set of balls. For a skin curve in $2D$ associated with two weighted points: see Figure 1.

2.2 Delaunay triangulation.

The Delaunay triangulation and Voronoi diagram are used to decompose the skin surface into patches of spheres and hyperboloids. We briefly give the definition of these structures and mention some properties.

The (weighted) Voronoi diagram (or: the power diagram) $\mathrm{Vor}(\mathcal{P})$ of a set of weighted points $\mathcal{P}$ is the subdivision of $\mathbb{R}^d$ into cells $\nu_{\mathcal{X}}$ that have smaller power distance to the weighted points in $\mathcal{X} \subseteq \mathcal{P}$ than to any other weighted point in $\mathcal{P}$:

$$\nu_{\mathcal{X}} = \bigcap_{\hat{p}\in\mathcal{X},\hat{p}'\in\mathcal{P}} \{x \in \mathbb{R}^d | \pi(\hat{p}, x) \leq \pi(\hat{p}', x)\}.$$

Write $\langle u, v\rangle$ for the inner product of u and v.

Observation 1. *Let $y_{\hat{p},\hat{p}'}$ be a point with the same power distance to $\hat{p}$ and $\hat{p}'$, then $\nu_{\mathcal{X}} = \bigcap_{\hat{p}\in\mathcal{X},\hat{p}'\in\mathcal{P}}\{x \in \mathbb{R}^d | \langle x - y_{\hat{p},\hat{p}'}, p' - p\rangle \leq 0\}$*

The dual of the Voronoi diagram is the Delaunay triangulation (or: regular triangulation) $\mathrm{Del}(\mathcal{P})$. We denote a Delaunay simplex of a set $\mathcal{X} \subseteq \mathcal{P}$, with $\nu_{\mathcal{X}} \neq \emptyset$, by $\delta_{\mathcal{X}}$. Recall that $\delta_{\mathcal{X}} = \texttt{conv}\,(\{p|\hat{p} \in \mathcal{X}\})$. If $\mathcal{X} \subsetneq \mathcal{X}'$ and $\nu_{\mathcal{X}'} \neq \emptyset$, then $\nu_{\mathcal{X}'}$ is a proper face of $\nu_{\mathcal{X}}$ and $\delta_{\mathcal{X}}$ is a proper face of $\delta_{\mathcal{X}'}$.

Observation 2. *The affine hulls of a Delaunay simplex $\delta_{\mathcal{X}}$ and its dual Voronoi cell $\nu_{\mathcal{X}}$ are complementary and orthogonal.*

Hence, the affine hulls of $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ always intersect in a single point, the center $f(\mathcal{X})$ of $\mathcal{X}$.

General position. In the remainder of this paper we assume general position, by which we mean that no $d+2$ weighted points are equidistant to a point in $\mathbb{R}^d$ and no $k+2$ centers of weighted points lie on a common k-flat for $k = 0, \ldots, d-1$. Several methods like [15] exist to symbolically perturb a data set and ensure these conditions. Note that, under this genericity condition, an orthosphere of a set $\mathcal{X}$ only exists if $|\mathcal{X}| \leq d+1$.

Consider a Delaunay cell $\delta_{\mathcal{X}'}$ and one of its faces $\delta_{\mathcal{X}}$, with $\mathcal{X} \subsetneq \mathcal{X}'$. Their duals are respectively a face of a Voronoi cell and the Voronoi cell itself. There is a half space through $\delta_{\mathcal{X}}$ containing $\delta_{\mathcal{X}'}$ and a half space through $\nu_{\mathcal{X}'}$ containing $\nu_{\mathcal{X}}$ such that their normals point in opposite directions.

Lemma 3. *Let $\delta_{\mathcal{X}}, \delta_{\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, such that $\delta_{\mathcal{X}}$ is a proper face of $\delta_{\mathcal{X}'}$ and let $u = x'_d - x_d$ with $x_d \in \delta_{\mathcal{X}}$, $x'_d \in \texttt{int}(\delta_{\mathcal{X}'})$. Then*

1. *$\langle u, x_v - x'_v\rangle = 0$, for $x_v, x'_v \in \nu_{\mathcal{X}'}$.*
2. *$\langle u, x_v - x'_v\rangle < 0$, for $x_v \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$, $x'_v \in \nu_{\mathcal{X}'}$*

Proof. Claim (1) follows directly from Observation 2. Hence, claim (2) is independent of the choice of x'_v.

For the proof of claim (2), let $m = |\mathcal{X}|$, $n = |\mathcal{X}'|$ and $\mathcal{X}' = \{\hat{p}_1, \ldots, \hat{p}_n\}$, such that $\hat{p}_i \in \mathcal{X}$, for $i \leq m$. Write x_d and x'_d in barycentric coordinates:

$x_d = \sum \gamma_i \cdot p_i$, $x'_d = \sum \gamma'_i \cdot p_i$ with $\sum \gamma_i = \sum \gamma'_i = 1$, $\gamma_i, \gamma'_i \geq 0$. Since $x_d \in \delta_{\mathcal{X}}$, $\gamma_i = 0$ for $i \geq m+1$, and $\gamma'_i > 0$ since $x'_d \in \texttt{int}(\delta_{\mathcal{X}'})$. Rewrite u as:

$$\begin{aligned} u &= x'_d - x_d \\ &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i) p_i + \sum_{i=m+1}^{n} \gamma'_i p_i \\ &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i)(p_i - p_1) + \sum_{i=m+1}^{n} \gamma'_i (p_i - p_1) \end{aligned}$$

Expanding $\langle u, x_v - x'_v \rangle$ yields:

$$\begin{aligned} \langle u, x_v - x'_v \rangle &= \sum_{i=1}^{m} (\gamma'_i - \gamma_i) \langle p_i - p_1, x_v - x'_v \rangle \\ &\quad + \sum_{i=m+1}^{n} \gamma'_i \langle p_i - p_1, x_v - x'_v \rangle \end{aligned}$$

From Observation 1, with $y_{\hat{p}_i, \hat{p}_1} = x'_v$, it follows that $\langle p_i - p_1, x_v - x'_v \rangle$ is not positive. Moreover, by Observation 2, the inner product is zero if and only if $\hat{p}_i \in \mathcal{X}$. Hence, the elements of the first sum are zero and the elements of the second sum are negative, so $\langle u, x_v - x'_v \rangle < 0$. Note that for $\hat{p}_1$ we can substitute any weighted point in $\mathcal{X}$. □

Although the meshing algorithm generalizes to any dimension, the main application is in $\mathbb{R}^3$. Therefore we present the algorithm in three space.

2.3 The mixed complex

The mixed complex $\text{Mix}^s(\mathcal{P})$, associated with a scalar $s \in [0, 1]$, is an intermediate complex between the Delaunay triangulation and the Voronoi diagram. Each mixed cell in the mixed complex is obtained by taking Minkowski sums of shrunken Delaunay simplices and their dual Voronoi cells.

Definition 4. *For $\delta_{\mathcal{X}} \in \text{Del}(\mathcal{P})$ the* mixed cell $\mu^s_{\mathcal{X}}$ *is defined by $\mu^s_{\mathcal{X}} = (1-s) \cdot \delta_{\mathcal{X}} \oplus s \cdot \nu_{\mathcal{X}}$.*

Here $\cdot$ denotes the multiplication of a set by a scalar and $\oplus$ denotes the Minkowski sum. For $s = 0$ the mixed cell is the Delaunay cell. When s increases it deforms affinely into the Voronoi cell for $s = 1$.

Each mixed cell is a convex polyhedron since it is the Minkowski sum of two convex polyhedra. Based on the dimension of the Delaunay simplex, there are four types of mixed cells. A mixed cell of type ℓ corresponds to a Delaunay ℓ-cell and is of the form $\mu^s_{\mathcal{X}}$ with $|\mathcal{X}| = \ell + 1$. In $3D$, mixed cells of type 3 are

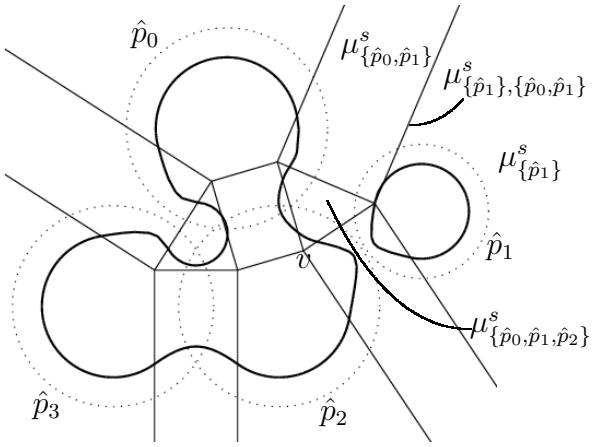

Figure 2: The skin curve of four weighted points (the dotted circles). Each mixed cell contains parts of an hyperbola or a circle. Some labels of mixed cells are given. Note that $v = \mu^s_{\{\hat{p}_0,\hat{p}_2\},\{\hat{p}_1,\hat{p}_2\}} = \mu^s_{\{\hat{p}_2\},\{\hat{p}_0,\hat{p}_1,\hat{p}_2\}}$.

tetrahedra (shrunken Delaunay 3-cells) and mixed cells of type 0 are shrunken Voronoi 3-cells. A mixed cell of type 1 or 2 is a prism with respectively the shrunken Voronoi facet or the shrunken Delaunay facet as its base.

The intersection of the skin surface and a mixed cell is a piece of a sphere or a hyperboloid. In the plane, the intersection of a skin curve with a mixed cell is either part of a circle or hyperbola. An example of the mixed complex and a skin curve is given in Figure 2. All rectangles are mixed cells of type 1 and contain hyperbolic patches. The other cells contain circular arcs. Depending on whether the mixed cell is of type 0 or 2, the interior of the skin curve lies inside or outside the circle.

Within a mixed ℓ-cell $\mu^s_{\mathcal{X}}$, the skin surface is a quadratic surface of the form $I_{\mathcal{X}}^{-1}(0)$, with:

$$I_{\mathcal{X}}(x) = -\frac{1}{1-s}\sum_{i=1}^{\ell} x_i^2 + \frac{1}{s}\sum_{i=\ell+1}^{3} x_i^2 - R^2, \tag{4}$$

and $x = (x_1, x_2, x_3)$. More precisely, $\mathtt{skn}^s\mathcal{X} \cap \mu^s_{\mathcal{X}} = I_{\mathcal{X}}^{-1}(0) \cap \mu^s_{\mathcal{X}}$. The coordinate system is orthonormal with its origin at the center of $\mathcal{X}$, and such that the first ℓ coordinates span the affine hull $\delta_{\mathcal{X}}$, see [14].

The following observation holds trivially for mixed cells of type 0 and 3. For mixed cells of type 1 and 2 it follows from the choice of the coordinate axis and the construction of the mixed cells.

Observation 5. *Each proper face of a mixed cell $\mu^s_{\mathcal{X}}$ is perpendicular to a symmetry set of $I_{\mathcal{X}}$*

Since the symmetry axis and the symmetry plane of the hyperboloid are

perpendicular, each face of a mixed cell of type 1 or 2 is parallel to the other symmetry set.

Polyhedral complex. The mixed complex is a polyhedral complex. The 3-cells of this polyhedral complex are formed by the mixed cells. We give a more detailed description of its structure.

Definition 6. *For $\mathcal{X}, \mathcal{X}' \in \mathcal{P}$, with $\nu_{\mathcal{X}}, \nu_{\mathcal{X}'} \neq \emptyset$, a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is defined as $\mu^s_{\mathcal{X},\mathcal{X}'} = \mu^s_{\mathcal{X}} \cap \mu^s_{\mathcal{X}'}$.*

Edelsbrunner gives an intuitive picture of the mixed complex in [14]. Take the interval of d-dimensional affine subspaces of $\mathbb{R}^{d+1}$ defined by $x_{d+1} = s$, for $s \in [0, 1]$. Draw $\mathrm{Del}(P)$ in $x_{d+1} = 0$ and $\mathrm{Vor}(\mathcal{P})$ in $x_{d+1} = 1$. For each Delaunay simplex and corresponding Voronoi cell construct

$$\mu_{\mathcal{X}} = \texttt{conv}\ (\delta_{\mathcal{X}} \cup \nu_{\mathcal{X}}).$$

All $\mu_{\mathcal{X}}$ are convex polyhedra of dimension $d+1$, their interiors mutually disjoint, and they decompose the strip between $x_{d+1} = 0$ and $x_{d+1} = 1$. The subspace $x_{d+1} = s$ intersects $\mu_{\mathcal{X}}$ in the mixed cell $\mu^s_{\mathcal{X}}$.

It is clear that a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is non-empty, for $0 < s < 1$, if the Delaunay and Voronoi cells of $\mathcal{X}$ and $\mathcal{X}'$ have a non-empty intersection. Or, equivalently, if $\nu_{\mathcal{X}\cap\mathcal{X}'}, \nu_{\mathcal{X}\cup\mathcal{X}'} \in \mathrm{Vor}(\mathcal{P})$. It is not enough for one of the two simplices to exists. E.g., let $\hat{p}_1, \hat{p}_2, \hat{p}_3$ be weighted points, the centers of which are the vertices of a triangle in a two-dimensional Delaunay triangulation and $\mathcal{X} = \{\hat{p}_1, \hat{p}_2\}$, $\mathcal{X}' = \{\hat{p}_3\}$. Then $\delta_{\mathcal{X}\cup\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, but $\mathcal{X} \cap \mathcal{X}' = \emptyset$, hence $\delta_{\mathcal{X}\cap\mathcal{X}'} \notin \mathrm{Del}(\mathcal{P})$. On the other hand, let $\delta_{\mathcal{X}}$, $\delta_{\mathcal{X}'}$ be two Delaunay edges, that share a common vertex, but do not have an incident triangle in common, then $\delta_{\mathcal{X}\cap\mathcal{X}'} \in \mathrm{Del}(\mathcal{P})$, but $\delta_{\mathcal{X}\cup\mathcal{X}'} \notin \mathrm{Del}(\mathcal{P})$.

For nonempty polyhedral cells, the following lemma describes the structure of the mixed complex.

Lemma 7. *A mixed cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is not empty iff $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are nonempty In that case,*

$$\mu^s_{\mathcal{X},\mathcal{X}'} = (1-s) \cdot \delta_{\mathcal{X}\cap\mathcal{X}'} \oplus s \cdot \nu_{\mathcal{X}\cup\mathcal{X}'}.$$

Before we prove this lemma we first make some general remarks. Note that the lemma also holds if $\nu_{\mathcal{X}\cap\mathcal{X}'} = \emptyset$ or $\nu_{\mathcal{X}\cup\mathcal{X}'} = \emptyset$.

Corollary 8. *If $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are nonempty, then $\mu^s_{\mathcal{X},\mathcal{X}'} = \mu^s_{\mathcal{X}\cap\mathcal{X}',\mathcal{X}\cup\mathcal{X}'}$*

The corollary holds since, $\mu^s_{\mathcal{X}\cap\mathcal{X}',\mathcal{X}\cup\mathcal{X}'} = (1-s) \cdot \delta_{\mathcal{X}\cap\mathcal{X}'} \oplus s \cdot \nu_{\mathcal{X}\cup\mathcal{X}'}$. Hence, each polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ has a unique label, if $\mathcal{X} \subset \mathcal{X}'$. To gain some intuition for the lemma, take s equal to zero. Then the mixed complex is the Delaunay triangulation and indeed $\mu^0_{\mathcal{X},\mathcal{X}'} = \delta_{\mathcal{X}} \cap \delta_{\mathcal{X}'} = \delta_{\mathcal{X}\cap\mathcal{X}'}$. Conversely, for $s = 1$, the mixed complex is the Voronoi diagram and $\mu^1_{\mathcal{X},\mathcal{X}'} = \nu_{\mathcal{X}} \cap \nu_{\mathcal{X}'} = \nu_{\mathcal{X}\cup\mathcal{X}'}$.

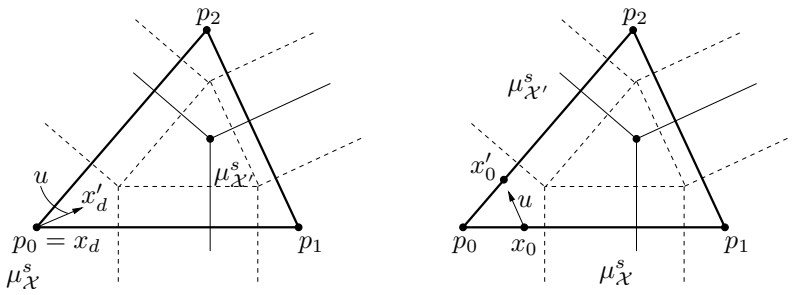

(a) $\mathcal{X} = \{\hat{p}_0\}, \mathcal{X}' = \{\hat{p}_0, \hat{p}_1, \hat{p}_2\}$, hence $\mathcal{X} \subseteq \mathcal{X}'$

(b) $\mathcal{X} = \{\hat{p}_0, \hat{p}_1\}, \mathcal{X}' = \{\hat{p}_0, \hat{p}_2\}$, hence $\mathcal{X} \not\subset \mathcal{X}'$ and $\mathcal{X}' \not\subset \mathcal{X}$

Figure 3: Illustration of the proof of Lemma 7.

For the proof of Lemma 7, we use a relation between the Delaunay cells and Voronoi cells, which we prove first. Consider a Delaunay cell $\delta_{\mathcal{X}'}$ and one of its faces $\delta_{\mathcal{X}}$, with $\mathcal{X} \subsetneq \mathcal{X}'$. Their duals are respectively a face of a Voronoi cell and the Voronoi cell itself. We show that there is a half space through $\delta_{\mathcal{X}}$ containing $\delta_{\mathcal{X}'}$ and a half space through $\nu_{\mathcal{X}'}$ containing $\nu_{\mathcal{X}}$ such that their normals point in opposite directions.

Proof of Lemma 7. The proof is trivial if $\mathcal{X} = \mathcal{X}'$, hence we assume that $\mathcal{X} \neq \mathcal{X}'$. For simplicity, let $F = (1-s) \cdot \delta_{\mathcal{X} \cap \mathcal{X}'} \oplus s \cdot \nu_{\mathcal{X} \cup \mathcal{X}'}$.

From Definition (4) it follows that $F \subseteq \mu^s_{\mathcal{X}}$ and $F \subseteq \mu^s_{\mathcal{X}'}$, since $\delta_{\mathcal{X} \cap \mathcal{X}'} \subseteq \delta_{\mathcal{X}}, \delta_{\mathcal{X}'}$ and $\nu_{\mathcal{X} \cup \mathcal{X}'} \subseteq \nu_{\mathcal{X}}, \nu_{\mathcal{X}'}$. Hence, $F \subseteq \mu^s_{\mathcal{X},\mathcal{X}'}$.

For the opposite inclusion, we show that the two mixed cells lie in opposite half spaces and intersect the bounding plane in F.

We distinguish two cases. First, consider the case where $\mathcal{X} \subseteq \mathcal{X}'$ or $\mathcal{X}' \subseteq \mathcal{X}$; See Figure 3(a). Without loss of generality we assume that $\mathcal{X} \subseteq \mathcal{X}'$. Let u be a vector perpendicular to $\delta_{\mathcal{X}}$ pointing from a point in $\delta_{\mathcal{X}}$ towards a point in the interior of $\delta_{\mathcal{X}'}$, such that $\langle u, x'_d - x_d \rangle > 0$, for $x_d \in \delta_{\mathcal{X}}$, $x'_d \in \delta_{\mathcal{X}'} \backslash \delta_{\mathcal{X}}$. Such a vector u exists, since $\delta_{\mathcal{X}}$ is a proper face of the convex polyhedron $\delta_{\mathcal{X}'}$. Note that u is perpendicular to $\delta_{\mathcal{X}}$. Lemma 3(2) states that $\langle u, x'_v - x_v \rangle < 0$, for $x_v \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$, $x_v \in \nu_{\mathcal{X}'}$.

For each point x in a mixed cell $\mu^s_{\mathcal{X}}$ there exists a unique combination $x_d \in \delta_{\mathcal{X}}$, $x_v \in \nu_{\mathcal{X}}$, such that $x = (1-s) \cdot \delta_{\mathcal{X}} + s \cdot \nu_{\mathcal{X}}$, since $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ are affinely independent. Hence, since $F \subseteq \mathcal{X}$, a point $y_0 \in F$ can be uniquely written as $y_0 = (1-s) \cdot y^0_d + s \cdot y^0_v$ with $y^0_d \in \delta_{\mathcal{X}}$, $y^0_v \in \nu_{\mathcal{X}'}$.

We analyse the sign of the inner product $\langle u, y - y_0 \rangle$ for y subsequently in $\mu^s_{\mathcal{X}} \setminus F$, F and $\mu^s_{\mathcal{X}'} \setminus F$.

First, let $y \in \mu^s_{\mathcal{X}} \setminus F$. We write $y = (1-s) \cdot y_d + s \cdot y_v$, with $y_d \in \delta_{\mathcal{X}}$ and $y_v \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X}'}$. The inner product $\langle u, y_d - y^0_d \rangle$ is zero since $y_d, y^0_d \in \delta_{\mathcal{X}}$ and $\langle u, y_v - y^0_v \rangle < 0$ by Lemma 3(2). Hence $\langle u, y - y_0 \rangle < 0$ for $y \in \mu^s_{\mathcal{X}} \setminus F$.

Now assume that $y \in F$. Similar to y_0, we write $y = (1-s) \cdot y_d + s \cdot y_v$, with

$y_d \in \delta_{\mathcal{X}}$, $y_v \in \nu_{\mathcal{X}'}$. The inner product $\langle u, y_d - y_d^0 \rangle$ is zero since $y_d, y_d^0 \in \delta_{\mathcal{X}}$ and $\langle u, y_v - y_v^0 \rangle = 0$ by Lemma 3*(1)*.

Finally, assume that $y \in \mu^s_{\mathcal{X}'} \setminus F$, then we write y as $y = (1-s) \cdot y_d + s \cdot y_v$, with $y_d \in \delta_{\mathcal{X}'} \setminus \delta_{\mathcal{X}}$ and $y_v \in \nu_{\mathcal{X}'}$. The inner product $\langle u, y_d - y_d^0 \rangle$ is positive by construction of u, and $\langle u, y_v - y_v^0 \rangle = 0$, again by Lemma 3*(1)*. Hence $\langle u, y - y_0 \rangle > 0$ for $y \in \mu^s_{\mathcal{X}'} \setminus F$.

Summarizing, we have:

$$\langle u, y - y_0 \rangle \begin{cases} < 0, \text{ for } y \in \mu^s_{\mathcal{X}} \setminus F, \\ = 0, \text{ for } y \in F, \\ > 0, \text{ for } y \in \mu^s_{\mathcal{X}'} \setminus F. \end{cases}$$

Hence, $\mu^s_{\mathcal{X}}$ and $\mu^s_{\mathcal{X}'}$ lie in opposite half spaces and meet only in F.

Conversely, assume that $\delta_{\mathcal{X}}$ is not a face of $\delta_{\mathcal{X}'}$ and vice versa. Then $\mathcal{X} \cap \mathcal{X}' \subsetneq \mathcal{X}, \mathcal{X}' \subsetneq \mathcal{X} \cup \mathcal{X}'$, viz. Figure 3(b). For this case the proof is similar, except for the construction of the vector u.

Let $x_I \in \delta_{\mathcal{X} \cap \mathcal{X}'}$. The Delaunay simplex $\delta_{\mathcal{X} \cap \mathcal{X}'}$ has at least co-dimension 2, since $|\mathcal{X} \cup \mathcal{X}'| - |\mathcal{X} \cap \mathcal{X}'| \geq 2$. Hence, the set of points orthogonal to $\delta_{\mathcal{X} \cap \mathcal{X}'}$ through x_I is at least 2-dimensional. We intersect this orthogonal set with a small ball centered at x_I. If the radius is small enough, the intersection contains a point $x_0 \in \texttt{int}(\delta_{\mathcal{X}})$ and $x_0' \in \texttt{int}(\delta_{\mathcal{X}'})$.

Let $u = \gamma \cdot (x_0' - x_0)$, for some $0 < \gamma < 1$. By construction the triangle x_I, x_0, x_0' is perpendicular to $\delta_{\mathcal{X} \cap \mathcal{X}'}$. Since $\| x_I - x_0 \| = \| x_I - x_0' \|$, the triangle x_I, x_0, x_0' is an isosceles triangle. Hence, the angles $\angle x_I, x_0, x_0'$, $\angle x_I, x_0', x_0$ are equal and acute. As a result, for $y_0 \in \delta_{\mathcal{X} \cap \mathcal{X}'}$ we have:

$$\langle u, y - y_0 \rangle \begin{cases} < 0, \text{ for } y \in \delta_{\mathcal{X}} \setminus \delta_{\mathcal{X} \cap \mathcal{X}'}, \\ = 0, \text{ for } y \in \delta_{\mathcal{X} \cap \mathcal{X}'}, \\ > 0, \text{ for } y \in \delta_{\mathcal{X}'} \setminus \delta_{\mathcal{X} \cap \mathcal{X}'}. \end{cases}$$

Note that u points from x_0 towards the interior of $\delta_{\mathcal{X} \cup \mathcal{X}'}$. Hence, u satisfies Lemma 3 with respect to $\delta_{\mathcal{X}}$ and $\delta_{\mathcal{X} \cup \mathcal{X}'}$. Using a similar argument, $-u$ satisfies Lemma 3 with respect to $\delta_{\mathcal{X}'}$ and $\delta_{\mathcal{X} \cup \mathcal{X}'}$. So, for $y_0 \in \nu_{\mathcal{X} \cup \mathcal{X}'}$ we have:

$$\langle u, y - y_0 \rangle \begin{cases} < 0, \text{ for } y \in \nu_{\mathcal{X}} \setminus \nu_{\mathcal{X} \cup \mathcal{X}'}, \\ = 0, \text{ for } y \in \nu_{\mathcal{X} \cup \mathcal{X}'}, \\ > 0, \text{ for } y \in \nu_{\mathcal{X}'} \setminus \nu_{\mathcal{X} \cup \mathcal{X}'}. \end{cases}$$

Now we combine the results for the Delaunay simplices and the Voronoi cells to a statement for the mixed cell. Let $y_0 \in F$ and write y_0 as $y_0 = (1-s) \cdot y_d^0 + s \cdot y_v^0$ with $y_d^0 \in \delta_{\mathcal{X}}$, $y_v^0 \in \nu_{\mathcal{X}'}$.

Write $y \in \mu^s_{\mathcal{X}} \setminus F$ uniquely as $y = (1-s) \cdot y_d + s \cdot y_v$, with $y_d \in \delta_{\mathcal{X}}$ and $y_v \in \nu_{\mathcal{X}}$. Since $y \notin F$, either $y_d \notin \delta_{\mathcal{X} \cap \mathcal{X}'}$ or $y_v \notin \nu_{\mathcal{X} \cup \mathcal{X}'}$. Expand the inner product $\langle u, y - y_0 \rangle$ to $(1-s) \cdot \langle u, y_d - y_d^0 \rangle + s \cdot \langle u, y_v - y_v^0 \rangle$. Using the estimates above, we obtain $\langle u, y - y_0 \rangle < 0$

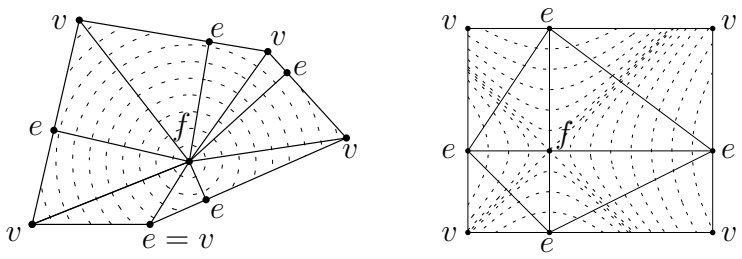

(a) Circular facets (b) Hyperbolic facets

Figure 4: The anchor points of two-dimensional polyhedral cells. Each anchor point is labeled by the type its cell (f for face, e for edge and v for vertex). The triangulation constructed in Section 3.2 is also shown.

A similar reasoning yields $\langle u, y - y_0 \rangle = 0$ for $y \in F$ and $\langle u, y - y_0 \rangle > 0$ for $y \in \mu^s_{\mathcal{X}} \setminus F$. □

Denote with $\mathtt{aff}(X)$ the affine hull of a set X.

Lemma 9. *For $\mathcal{X} \subset \mathcal{X}'$, $\mathtt{aff}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and $\mathtt{aff}(\delta_{\mathcal{X}'}) \cap \mathtt{aff}(\nu_{\mathcal{X}})$ are complementary and orthogonal.*

Proof. To shorten notation we write $F = \mathtt{aff}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and $G = \mathtt{aff}(\delta_{\mathcal{X}'}) \cap \mathtt{aff}(\nu_{\mathcal{X}})$. Recall from Lemma 7 that $\mu^s_{\mathcal{X},\mathcal{X}'} = (1-s) \cdot \delta_{\mathcal{X}} \oplus s \cdot \nu_{\mathcal{X}'}$.

The cells $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}'}$ are affinely independent, hence $\dim F = \dim \delta_{\mathcal{X}} + \dim \nu_{\mathcal{X}'} = d + |\mathcal{X}| - |\mathcal{X}'|$. Further, $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}}$ are orthogonal and $\dim G = \dim \mathtt{aff}(\delta_{\mathcal{X}'}) - \dim \mathtt{aff}(\delta_{\mathcal{X}}) = |\mathcal{X}'| - |\mathcal{X}|$. Hence, the dimensions of F and G add up to d. Both $\delta_{\mathcal{X}}$ and $\nu_{\mathcal{X}'}$ are orthogonal to G, which shows the orthogonality of F and G. □

Corollary 10. *The dimension of a non-empty mixed cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ in $\mathbb{R}^d$, is $d - |\mathcal{X} \cup \mathcal{X}'| + |\mathcal{X} \cap \mathcal{X}'|$.*

2.4 The anchor point.

For the construction of the mesh we use the anchor point of a polyhedron.

Definition 11. *Let A be a convex set and p a point in $\mathbb{R}^3$. Then the* anchor point *$a_p(A)$ is the point in A closest to p.*

We are interested in the case where A is a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$, a Delaunay cell $\delta_{\mathcal{X}}$ or a Voronoi cell $\nu_{\mathcal{X}}$ and p is the center $f(\mathcal{X})$. In fact, we use the anchor points of the polyhedral cells as vertices of a tetrahedral complex that decomposes the skin surface into topological disks.

We distinguish two types of critical points on a mixed cell $\mu^s_{\mathcal{X}}$, interior critical points are critical points of $I_{\mathcal{X}}$ contained in the interior of $\mu^s_{\mathcal{X}}$ and boundary

critical points are critical points of $I_{\mathcal{X}}$ restricted to the boundary of $\mu^s_{\mathcal{X}}$. All critical points are anchor points of a face of the mixed cell, viz. Figure 4. However, not all anchor points are critical points, e.g. the point that is both the anchor point of a vertex and an edge in Figure 4(a).

Lemma 12. *A (boundary or regular) critical point of $I_{\mathcal{X}}$ on a polyhedral cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ is the anchor point of $\mu^s_{\mathcal{X},\mathcal{X}'}$ or the anchor point of one of its faces with respect to $f(\mathcal{X})$.*

Proof. The center $f(\mathcal{X})$ is the only critical point of the quadratic function $I_{\mathcal{X}}$. If $f(\mathcal{X})$ is contained in $\mu^s_{\mathcal{X},\mathcal{X}'}$, then it is the anchor point $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'})$.

It remains to show that all boundary critical points are also anchor points. By Observation 5, a face of $\mu^s_{\mathcal{X},\mathcal{X}'}$ is either parallel or perpendicular to the symmetry sets of $I_{\mathcal{X}}$. Hence, if $f(\mathcal{X})$ projects onto the facet, then the facet has a boundary critical point. By definition, this point is the anchor point of the facet with respect to $f(\mathcal{X})$. □

Lemma 13. $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = a_{f(\mathcal{X}')}(\mu^s_{\mathcal{X},\mathcal{X}'})$.

Proof. Both $f(\mathcal{X})$ and $f(\mathcal{X}')$ lie on $\mathtt{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'})$ and $\mathtt{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$. Hence they lie on $\mathtt{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'}) \cap \mathtt{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$, which is orthogonal to $\mu^s_{\mathcal{X},\mathcal{X}'}$ by Lemma 9. □

Lemma 14. $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$.

Proof. If A and B are orthogonal, then $a_{f(\mathcal{X})}(s\,A \oplus (1-s)\,B) = a_{f(\mathcal{X})}(s\,A) + a_{f(\mathcal{X})}((1-s)\,B)$. Therefore, since $\delta_{\mathcal{X}\cap\mathcal{X}'}$ and $\nu_{\mathcal{X}\cup\mathcal{X}'}$ are orthogonal, we have $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X})}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X})}(\nu_{\mathcal{X}\cup\mathcal{X}'})$.

Since $f(\mathcal{X}), f(\mathcal{X}\cap\mathcal{X}') \in \mathtt{aff}(\nu_{\mathcal{X}\cap\mathcal{X}'})$ and $\delta_{\mathcal{X}\cap\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}\cap\mathcal{X}'}$, we have $a_{f(\mathcal{X})}(\delta_{\mathcal{X}\cap\mathcal{X}'}) = a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'})$. Similarly, $f(\mathcal{X}), f(\mathcal{X}\cup\mathcal{X}') \in \mathtt{aff}(\delta_{\mathcal{X}\cup\mathcal{X}'})$ and $\delta_{\mathcal{X}\cup\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}\cup\mathcal{X}'}$, hence $a_{f(\mathcal{X})}(\nu_{\mathcal{X}\cup\mathcal{X}'}) = a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$. Concluding, $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'}) = (1-s)\cdot a_{f(\mathcal{X}\cap\mathcal{X}')}(\delta_{\mathcal{X}\cap\mathcal{X}'}) + s\cdot a_{f(\mathcal{X}\cup\mathcal{X}')}(\nu_{\mathcal{X}\cup\mathcal{X}'})$. □

Now that we have the decomposition of the anchor point of a polyhedral cell into the anchor point of Delaunay and Voronoi cells, we show that these anchor points are easily constructed.

Lemma 15. *The anchor point $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ lies in the interior of $\delta_{\mathcal{X}}$ or $a_{f(\mathcal{X})}(\delta_{\mathcal{X}}) = a_{f(\mathcal{X}')}(\delta_{\mathcal{X}'})$, where $\delta_{\mathcal{X}'}$ is a face of $\delta_{\mathcal{X}}$*

Proof. Assume that $f(\mathcal{X})$ is not contained in $\mathtt{int}(\delta_{\mathcal{X}})$, otherwise the proof is trivial. Since $\delta_{\mathcal{X}}$ is a convex polyhedron, the point closest to $f(\mathcal{X})$ lies on a proper face of $\delta_{\mathcal{X}}$, say $\delta_{\mathcal{X}'}$.

Since $f(\mathcal{X}), f(\mathcal{X}') \in \mathtt{aff}(\nu_{\mathcal{X}'})$ and $\delta_{\mathcal{X}'}$ is orthogonal to $\nu_{\mathcal{X}'}$, we have $a_{f(\mathcal{X})}(\delta_{\mathcal{X}'}) = a_{f(\mathcal{X}')}(\delta_{\mathcal{X}'})$. □

A similar lemma holds for Voronoi cells, for which we omit the proof.

Lemma 16. *The anchor point $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$ lies in the interior of $\nu_{\mathcal{X}}$ or $a_{f(\mathcal{X})}(\nu_{\mathcal{X}}) = a_{f(\mathcal{X}')}(\nu_{\mathcal{X}'})$, where $\nu_{\mathcal{X}'}$ is a face of $\nu_{\mathcal{X}}$*

Concluding, with the anchor points $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ and $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$, with $\delta_{\mathcal{X}} \in \mathrm{Del}(\mathcal{P})$ we can construct the anchor point of any polyhedral cell. Moreover, Lemma 15 and Lemma 16 give a recursive definition that makes it easy to compute $a_{f(\mathcal{X})}(\delta_{\mathcal{X}})$ and $a_{f(\mathcal{X})}(\nu_{\mathcal{X}})$.

3 The meshing algorithm

This section describes the construction of a tetrahedral complex for which the intersection of a cell with the skin surface is either empty or a topological disk. Moreover we show that the mesh extracted from this tetrahedral complex by the marching tetrahedra algorithm [24] is isotopic to the skin surface.

3.1 Monotonicity condition.

In Section 3.2 we give a detailed construction of the tetrahedral complex. For now, we only give the main condition imposed on the tetrahedral complex. First, we require that each tetrahedron is contained in a single mixed cell. Recall that the skin surface restricted to a mixed cell $\mu^s_{\mathcal{X}}$ is a subset of the quadric $I^{-1}_{\mathcal{X}}(0)$, cf. Equation (4). Express a point $x = (x_1, x_2, x_3)$ in the local coordinate system of $I_{\mathcal{X}}$.

Condition 17 (Monotonicity). *Let ab be a line segment contained in a mixed cell $\mu^s_{\mathcal{X}}$ of type ℓ, with $I_{\mathcal{X}}(a) \leq I_{\mathcal{X}}(b)$. The segment ab satisfies the monotonicity condition if $x_1^2 + \ldots + x_\ell^2$ is non-increasing and $x_{\ell+1}^2 + \ldots + x_3^2$ is non-decreasing on the segment from a to b.*

In words, a segment ab satisfies the monotonicity condition if the distance to both symmetry sets of $I_{\mathcal{X}}$ is monotone and the distance to one symmetry set of the quadric does not increase when the distance to the other symmetry set increases. For spheres ($\ell = 0, 3$) one symmetry set is empty and the monotonicity condition is satisfied if the distance to the center of the sphere is monotone. For hyperboloids ($\ell = 1, 2$) a segment satisfies the monotonicity condition if the distances to the symmetry axis and the symmetry plane are monotone and the distance to one symmetry set does not increase when the distance to the other symmetry set increases. From Equation (4) we conclude:

Observation 18. *If a line segment ab satisfies the monotonicity condition, then $I_{\mathcal{X}}$ is monotonically increasing on ab.*

We construct the tetrahedral complex in such a way that all edges satisfy the monotonicity condition. In fact, if all edges satisfy the monotonicity condition, then a generalized monotonicity condition holds for all cells.

Lemma 19. *Let $\mu^s_{\mathcal{X}}$ be a mixed cell of type ℓ and let $v_1, \ldots, v_n$ be the vertices of a cell of the tetrahedral complex in $\mu^s_{\mathcal{X}}$, with $I_{\mathcal{X}}(v_i) \leq I_{\mathcal{X}}(v_j)$ if $i < j$.*

If the monotonicity condition holds for all edges then, each segment ab, with $a \in \texttt{conv}\,(v_1, \ldots, v_k)$ and $b \in \texttt{conv}\,(v_{k+1}, \ldots, v_n)$, for $k \in \{1, \ldots, n\}$, satisfies the monotonicity condition.

In the proof we need a small lemma.

Lemma 20. *Let $v_1v_2v_3$ be a triangle in $\mathbb{R}^2$, such that the distance to the origin is monotonically increasing (decreasing) along both v_1v_3 and v_2v_3. Then the distance to the origin is monotonically increasing (decreasing) on the segment xv_3, with $x \in v_1v_2$.*

Proof. Let $x = (1-t)v_1 + tv_2$ and let $d(t) = \| (1-t)x + tv_3 \|^2$ be the squared distance to the origin on the line segment xv_3. The distance $d(t)$ is monotone if $d'(t) \neq 0$ for $t \in [0,1]$. Since $d'(t) = 2\langle x, v_3 - x\rangle + 2t\langle v_3 - x, v_3 - x\rangle$, $d(t)$ is monotone if $\langle x, v_3 - x\rangle \geq 0$ or $\langle v_3, v_3 - x\rangle \leq 0$.

Assume that the distance to the origin increases monotonically on both line segments v_1v_3 and v_2v_3, hence $\langle v_1, v_3 - v_1\rangle \geq 0$ and $\langle v_2, v_3 - v_2\rangle \geq 0$. We have

$$\begin{aligned} \langle x, v_3 - x\rangle &= (1-t)\langle v_1, v_3 - v_1\rangle \\ &\quad +t\langle v_2, v_3 - v_2\rangle \\ &\quad +t(1-t)\| v_2 - v_1 \|^2 \\ &\geq 0. \end{aligned}$$

Conversely assume that the distance to the origin decreases monotonically on both line segments v_1v_3 and v_2v_3, hence $\langle v_3, v_3 - v_1\rangle \leq 0$ and $\langle v_3, v_3 - v_2\rangle \leq 0$. Then we have $\langle v_3, v_3 - x\rangle = (1-t)\langle v_3, v_3 - v_1\rangle + t\langle v_3, v_3 - v_2\rangle \leq 0$. □

Proof of Lemma 19. We repeatedly move vertices along edges of the cell of the tetrahedral complex while maintaining the monotonicity condition. After the displacement of the vertices, the line segment ab lies on one of the edges.

Since $I_{\mathcal{X}}(v_i) \leq I_{\mathcal{X}}(v_j)$ for $i < j$ and by Equation (4), the distance to the symmetry set spanned by the first ℓ coordinate axis is decreasing and the distance to the other symmetry set is increasing. Assume for now that when we move the vertices v_i and v_{i+1} over the edge v_iv_{i+1}, that all edges in the new cell satisfy the monotonicity condition.

We now move v_1 to a and b to v_n. If $k = 1$, then a lies on v_1. Otherwise a we move v_{k-1} to v_k until a lies on the face $\texttt{conv}\,(v_1, \ldots, v_{k-1})$ and repeat this step for $k-1$. Similarly we move v_{k+1} to v_k until b lies on the face $\texttt{conv}\,(v_{k+1}, \ldots, v_n)$ and repeat the step for $k+1$ until $k = n-1$. Then ab is an edge of the new tetrahedron. Hence ab satisfies the monotonicity condition.

It remains to show that the new edges also satisfy the monotonicity condition. Therefore consider three vertices v_i, v_{i+1} and v_j. If $j < i$ $(j > i+1)$ then

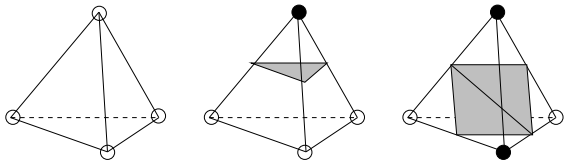

Figure 5: The three different configurations of a tetrahedron. White and black vertices lie on different sides of the skin surface.

the distance to the first symmetry set along the line segment v_jv_i and v_jv_{i+1} is decreasing (increasing) and to the second symmetry set it is increasing (decreasing). We show that the distances remain monotonically increasing or decreasing along a line segment xv_j for $x \in v_iv_{i+1}$. We distinguish three cases. First, assume that the symmetry set is a point. We project the symmetry point on the plane $v_jv_iv_{i+1}$. From the previous lemma it follows that the distance on xv_j to the projection of the symmetry point is monotone, and therefore also the distance to the symmetry point. Next, assume that the symmetry set is a line, then we project the triangle $v_jv_iv_{i+1}$ on a plane orthogonal to the symmetry line. By applying the previous lemma, it follows that the distance to the projection of the line on the plane along the line segment xv_j is monotone. Finally, if the symmetry set is a plane we do not need the previous lemma. In this case the distance to the plane at x is smaller (greater) than the distance at w_3. Hence the distance to the symmetry plane on xw_3 is monotone.

To conclude, if the distance to a symmetry set is monotonically increasing (decreasing) along both v_jv_i and v_jv_{i+1}, then this distance is also monotonically increasing (decreasing) along the line segment xv_j, with $x \in v_iv_{i+1}$. Hence xv_j satisfies the monotonicity condition. □

Mesh extraction. The coarse mesh is extracted from the tetrahedral complex by the marching tetrahedra algorithm [24]. Each edge of the tetrahedral complex intersects the skin surface at most once by Observation 18. We place vertices of the mesh on these intersection points. Then the mesh is constructed by considering the number of vertices of the tetrahedron inside the skin surface as depicted in Figure 5. The third configuration remains ambiguous, since the common interior edge of the two triangles can be flipped.

Theorem 21. *A tetrahedral complex for which each edge satisfies the monotonicity condition has two properties:*

1. *each cell intersects the skin surface in a topological disk and*
2. *the mesh extracted from the tetrahedral complex is isotopic to the skin surface.*

Proof. Let V^- and V^+ be the vertices of a k-cell of the tetrahedral complex inside and outside the skin surface, respectively. Consider the set of line segments ab with $a \in \texttt{conv}\,(V^-)$, $b \in \texttt{conv}\,(V^+)$. The set of line segments is empty if the cell does not intersect the skin surface, i.e. if $V^- = \emptyset$ or $V^+ = \emptyset$. On the other hand, if the cell intersects the skin surface, then the set of line segments spans the cell and the line segments may intersect but only at their endpoints. On faces of the cell, the line segments are defined consistently because there the construction is based only on the labels of vertices of the face.

By Lemma 19, each segment satisfies the monotonicity condition. Hence $I_{\mathcal{X}}$ is monotone on ab. Moreover, a lies inside and b outside the skin surface. Therefore ab intersects the skin surface in a single point. Since the segments span the tetrahedron, the skin surface within the cell is a topological disk.

A simple case study shows that each segment also intersects the coarse mesh transversally in exactly one point. The isotopy is constructed by parameterizing the mesh and the skin surface with these segments and using linear interpolation between the parameterizations. □

We call the segments in the proof above *transversal segments* because each segment intersects both the skin surface and the coarse mesh transversally in a single point.

3.2 The tetrahedral complex.

Up to now we assumed that it is possible to construct a tetrahedral complex in such a way that all edges satisfy the monotonicity condition. In this section we construct this tetrahedral complex. We triangulate polyhedral cells in order of increasing dimension.

All vertices of the tetrahedral complex are anchor points of polyhedral cells. In case an anchor point lies on the boundary of its polyhedral cell, it coincides with another anchor point and the simplicial complex is degenerate. Therefore, during the construction of the tetrahedral complex we test whether the anchor point lies in the interior of the polyhedral cell, and collapse the vertex otherwise. For simplicity, in the remainder of this section we assume that the anchor point lies in the interior of the mixed cell.

Subdividing polyhedral cells of positive co-dimension. On each vertex of the polyhedral complex we place a vertex of the tetrahedral complex. Note that these vertices are the anchor point of 0-cells of the polyhedral complex.

Next, consider an edge $\mu^s_{\mathcal{X},\mathcal{X}'}$ of the polyhedral complex. By Lemma 12, if $I_{\mathcal{X}}$ has a critical point on the interior of the edge, this critical point is the anhor point of the edge. Therefore we split the edge in the anchor point $a_{f(\mathcal{X})}(\mu^s_{\mathcal{X},\mathcal{X}'})$ and construct two edges from the anchor point to the vertices. By Observation 5

a polyhedral edge is parallel to one symmetry set and is split in the point closest to the projection of the other symmetry set, hence both edges satisfy the monotonicity condition.

We distinguish two types of facets: *circular* facets are facets for which the contour lines of $I_{\mathcal{X}}$ restricted to the facet are circles. The other facets are called *hyperbolic* because the contour lines are hyperbolas on the facet. Since the skin surface is tangent continuous, $I_{\mathcal{X}}|\mu^s_{\mathcal{X},\mathcal{X}'} = I_{\mathcal{X}'}|\mu^s_{\mathcal{X},\mathcal{X}'}$ and the facet inherits the same type from both mixed cells $\mu^s_{\mathcal{X}}$ and $\mu^s_{\mathcal{X}'}$. All facets of mixed cells of type 0 and 3 are spherical. The facets of a mixed cell of type 1 or 2 are spherical if they touch a mixed cell of type 0 or 3, and hyperbolic if they touch a mixed cell of type 1 or 2.

We triangulate circular and hyperbolic facets differently. Circular facets are triangulated by adding an edge from the anchor point of the facet to each anchor point on the boundary of the facet, i.e., either the anchor point of an edge or a vertex. See Figure 4(a). Since the anchor point of the facet is the point closest to the center of the sphere, the distance to the center increases monotonically on each edge and each edge satisfies the monotonicity condition.

Hyperbolic facets are rectangles with edges parallel or perpendicular to the symmetry axis of the corresponding hyperboloid. The anchor point of an edge is the point closest to the center, hence it is the point on the edge closest to the symmetry axis the edge is orthogonal to. Similarly, the anchor point of the facet is the point on the facet closest to the center. Thus the edges from the anchor point of the facet to the anchor point of an edge are parallel to one axis and the distance to the other axis increases monotonically. Further, we add edges from the anchor point of an edge to the anchor point of an orthogonal edge. On these edges the distance to one symmetry axis increases whereas the distance to the other symmetry axis decreases. This triangulation is depicted in Figure 4(b).

Subdividing polyhedral cells of type 0 and 3. The mixed cells of this type contain a spherical patch of the skin surface. Similar to spherical facets, we have to triangulate polyhedral cells of type 0 and 3 in such a way that the distance to the center is monotone on each edge. The anchor point of the mixed cell is the point in the mixed cell closest to the center. Hence the distance to the center on each line segment from the anchor point of the mixed cell to any other point in the mixed cell is monotone, and therefore satisfies the monotonicity condition. We have already constructed the triangulation of the boundary of the mixed cell and triangulate the entire cell by adding edges from the anchor point of the cell to each vertex on the boundary. The tetrahedra are formed by taking the join of a triangle on the the triangulated boundary of the mixed cell and the anchor point of the mixed cell.

Subdividing polyhedral cells of type 1 and 2. The triangulation of mixed cells of type 1 and 2 is slightly more subtle. The mixed cell contains a hyperboloid patch of the skin surface and the mixed cell is a prism with its base parallel to the symmetry plane of the hyperboloid. For an edge to satisfy the monotonicity condition, the distance to both the symmetry plane and the symmetry axis has to be monotone and the distance to one symmetry set may not increase, when the distance to the other symmetry set increases.

We already triangulated the facets of the mixed cell. The hyperbolic facets of the prism are the facets that are parallel to the symmetry axis. We split the prism in the plane V through the anchor point of the mixed cell parallel to the symmetry plane. This plane also contains the anchor points of the faces and edges of the mixed cells that are parallel to the symmetry axis. Hence each facet parallel to the symmetry axis is already split in V. The new facet is spherical and we triangulate it accordingly.

Consider one split mixed cell. The base of the prism furthest away from the symmetry plane contains the points furthest away from the symmetry plane. Hence its anchor point is the point with maximal distance to the symmetry plane and minimal distance to the symmetry axis. Therefore, all line segments in the split mixed cell with this anchor point as a vertex satisfy the monotonicity condition. The boundary of the prism is already triangulated and we triangulate the split mixed cell by adding edges from the anchor point of the base to all vertices on the boundary. The tetrahedra are the join of a triangle on the triangulated boundary and the anchor point of the base.

Union of balls. For a shrink factor one, the skin surface of a set of balls is the union of these balls. In this case, the mixed complex is the Voronoi diagram. This means that only mixed cells of type 0 are three-dimensional cell. This greatly simplifies the set of tetrahedra.

It is also desirable to retain edges of the mesh on the intersection of two balls. The subdivision algorithm ensures this by definition. Chew's algorithm can also be extended to allow constrained edges, see [11].

3.3 Complexity analysis.

In many real world applications the size of the Delaunay triangulation is linear in the number of input balls, see [12, 2]. However, the worst case complexity of the Delaunay triangulation is quadratic in the number of input balls, see [7]. We show that the size of the coarse mesh is linear in the size of the Delaunay triangulation and that this is worst case optimal.

Lemma 22. *The size of the coarse mesh is linear in the size of the Delaunay triangulation.*

Because of the duality relationship, the size of the Voronoi complex is equal to the size of the Delaunay triangulation. The size of the mixed complex is linear in the size of the Delaunay triangulation, since the complexity of each mixed cell is linear in the complexity of the Delaunay and Voronoi cell it is constructed from. The number of tetrahedra within a mixed cell is linear in the complexity of the mixed cell and within each tetrahedron we construct at most two triangles. Thus, the mesh is linear in the size of the Delaunay triangulation.

To show that this is worst case optimal, we construct a skin surface with $O(n^2)$ holes from a set of n balls. Any mesh with $O(n^2)$ holes has complexity $O(n^2)$, thus giving the lower bound. The construction is done as follows: the first $n/2$ balls are centered on the unit circle in the xy-plane and have radius 0.5. The other $n/2$ balls are centered on the z-axis and their radius is such that they touch the first $n/2$ balls. The union of these balls has $O(n^2)$ holes because each ball on the z-axis adds $O(n)$ holes to the union. The skin surface also has $O(n^2)$ holes because it is isotopic to the union of the balls.

3.4 Mesh enhancement.

The topologically correct mesh obtained with the marching tetrahedra algorithm is rather coarse and may contain long and skinny triangles. Therefore, we develop a method to enhance the mesh while maintaining the isotopy. The changes to the mesh we allow are local and do not change the topology of the mesh.

Before we change the mesh, we first test whether the isotopy with the skin surface is maintained. Therefore we use the transversal segments as described in the proof of Theorem 21. In fact, we first test whether each transversal line segment intersects the new mesh exactly once. We conclude this section with two examples of mesh refinement algorithms.

Changing the mesh. To test whether the isotopy is maintained under a change of the mesh we would have to test whether each transversal line segment intersects the mesh once. We rephrase this in such a way that it is easier to verify.

Let c be a 3-cell in the tetrahedral complex, t a triangle of the mesh intersecting c and V^- and V^+ the vertices of c inside and outside the skin surface, respectively.

Lemma 23. *If for all t and c, V^- and V^+ are separated by the plane through t and V^- lies in the direction of the inner part of the mesh, then each transversal line segment within c intersects the mesh once.*

Proof. Consider a line segment p^-p^+, with $p\pm \in \texttt{conv}\,(V^\pm)$. Since p^- and p^+ lie on opposite sides of the mesh, p^-p^+ intersects the mesh at least once. Assume that it intersects p^-p^+ more than once, then on the second intersection point from p^-, the segment moves from outside the mesh to the inside. Hence. the inner product with the normal is negative. □

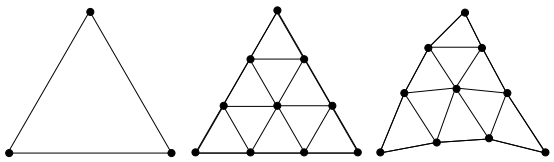

Figure 6: The Sqrt-3 subdivision method applied twice. Left: the original triangle, in the middle the subdivided triangle. On the right are the vertices placed on the skin surface.

We now have an efficient way of testing whether the isotopy of the skin surface and the mesh is maintained. If the test fails, then the mesh is too coarse and we refine the mesh. We show that the refinement succeeds for small triangles.

Lemma 24. *A triangle t of the mesh contained in a single tetrahedron c can be subdivided in any point $x \in t$ by moving x along the transversal segments to the skin surface.*

Proof. Since x moves along the line segments within c, the line segments through f and the subdivided faces are the same. Thus the new mesh can be obtained from the old mesh by interpolation along the line segments. □

We can also flip an edge of the triangulation if the two adjacent triangles and the new triangles are intersected by the same transversal segments. This condition is similar to the condition in a two-dimensional mesh that an edge can be flipped if the union of the two triangles is convex.

To summarize, we have an efficient test to check whether the isotopy is maintained. If a change of the mesh would result in a violation of the isotopy test, then we can always subdivide the face into faces that are contained within a single tetrahedron. Each of these faces satisfy the isotopy test.

Sqrt-3 method We implemented the sqrt-3 subdivision method [18] on the coarse mesh. The sqrt-3 subdivision method splits each triangle into 9 subtriangles and then moves the newly created vertices towards the skin surface along the transversal segments.

By Lemma 24, the subdivision algorithm maintains the isotopy. Hence, it is not necessary to test isotopy, which make the algorithm very fast. On the other hand, the subdivision algorithm does not improve the quality of the triangles. Therefore this method is not suitable for constructing a mesh for numerical simulations.

Chew's algorithm We also implemented Chew's algorithm [11] to improve the quality of the triangles of the coarse mesh and obtain a mesh suitable for numerical simulations. After the algorithm terminates, each triangle has angles

Molecule	Our algorithm			Dynamic	Marching
	Coarse	Sqrt-3	Chew		
pdb7tmn	0:00:01	0:00:02	0:00:05	0:10:00	0:00:05
DNA	0:00:14	0:00:29	0:00:55	0:35:12	0:00:51
Gramacidin A	0:00:08	0:00:31	0:01:13	1:35:23	0:03:22

Table 1: Performance comparison

between 30 and 120 degrees and has a user defined maximal size. The only constraint on the size-criterion is that there exists a $\delta > 0$ such that any well-shaped triangle that fits within a circle of radius δ is well-sized. We chose the size of a triangles inversely proportional to the maximal curvature which is nonzero on skin surfaces.

During the refinement, we test the isotopy before an inserting a new point and before flipping an edge.

4 Implementation.

We implemented the algorithm described above in C++ using CGAL [1]. First we compute the Delaunay triangulation of the weighted points. From this triangulation we extract the mixed complex and triangulate it, as described in Section 3.

We use filtering techniques [8] to increase the speed of the refinement algorithm. Therefore we compute the exact location of a anchor points (the vertices of the tetrahedral complex) and an interval containing the anchor point. We then use interval arithmetic to test the isotopy and revert to exact computations when the interval arithmetic is not exact.

We also implemented both algorithms described in the previous section. The implementation of the $\sqrt{3}$-subdivision method is straightforward. Before we apply Chew's algorithm we perform a preprocessing step in which we remove small edges. This reduces the size of the final mesh considerably.

5 Examples and experiments.

We compare our algorithm to the algorithms described in [9] and [10]. There is a comparison of the two algorithms in [10]. These tests are run on a Pentium 4 running at 2.54GHZ. To test our algorithm we used an AMD Athlon 1800+ which is actually a little slower. We tested our algorithm on various molecules, computing only the coarse mesh, the coarse mesh and one $\sqrt{3}$-subdivision step and the coarse mesh which subsequently improved using Chew's algorithm. For timings see Table 1.

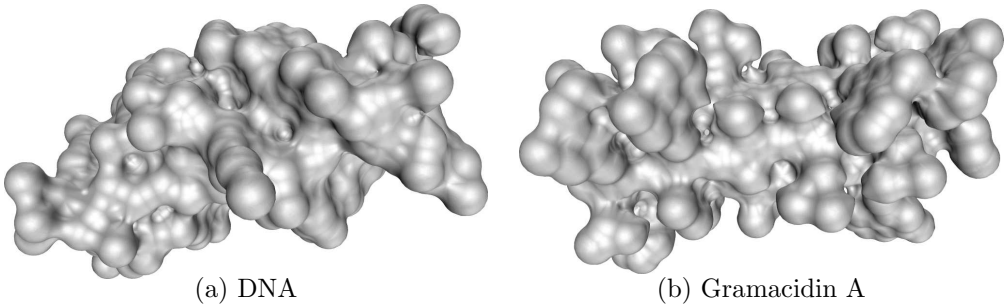

(a) DNA (b) Gramacidin A

Figure 7: Two larger molecules.

Note that both our algorithm and the marching algorithm [10] are significantly faster than the dynamic skin algorithm [9]. However, we believe that [10] does not come with topological guarantees.

Figure 8 shows the molecule pdb7tmn. In Figure 8(d) we enlarged a part of the coarse mesh and applied the Sqrt-3 method in Figure 8(e). Note that the triangles remain skinny. Figure 8(f) shows the result of applying Chew's algorithm directly to the coarse mesh. Because of small edges in the coarse mesh, there are also small edges near parts with low curvature. When we remove small edges, viz. Figure 8(g), before we apply Chew's algorithm, we obtain Figure 8(h).

6 Conclusion and future work.

We present an algorithm that constructs a mesh that is isotopic to the skin surface and discuss two methods to refine this mesh.

The algorithm we present is static in the sense that it generates a mesh for a fixed set of input balls. The rigid foundation of the tetrahedral complex makes us believe that it is also possible to maintain the coarse mesh while deforming the input set. This is important for deforming molecules. Two deformation schemes seem computationally interesting. From Lemma 14 we know that the anchor point of a mixed cell $\mu^s_{\mathcal{X},\mathcal{X}'}$ only depends on the Delaunay cell $\delta_{\mathcal{X}\cup\mathcal{X}'}$, the Voronoi cell $\nu_{\mathcal{X}\cap\mathcal{X}'}$ and the shrink factor. Adding a constant to all weights does not change the Delaunay and Voronoi diagram and hence does not change the simplicial complex. Hence, the coarse mesh of the skin surface obtained by adding a constant to each weight is another level-set of the tetrahedral complex. Another deformation is obtained by varying the shrink factor. Again, the structure of the simplicial complex remains unchanged, however the positions of the anchor points change. It is sufficient to reposition the anchor points and then update or recompute the coarse mesh.

We also believe that a similar algorithm can compute Connolly surfaces. These surfaces are also used in molecular biology and are formed by a small probe

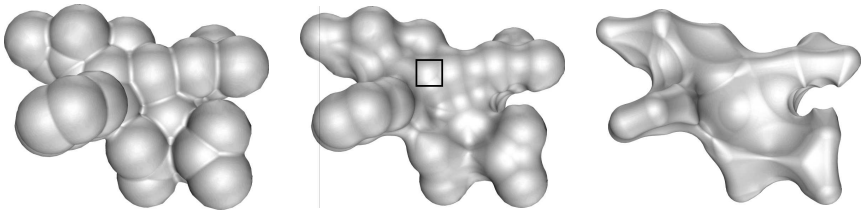

(a) Shrink: .85 (b) Shrink: .5 (c) Shrink: .15

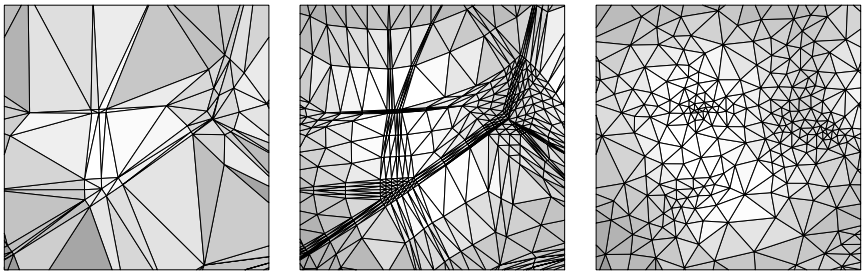

(d) Coarse mesh (e) Sqrt3-method (f) Chew applied to the coarse mesh

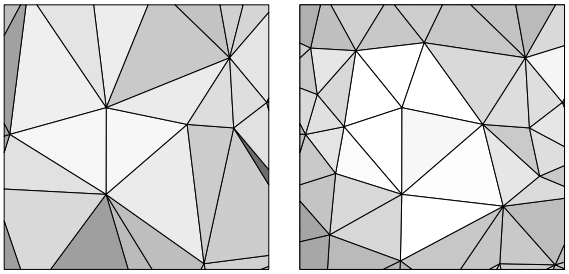

(g) Small edges removed (h) Chew applied to the enhanced mesh

Figure 8: pdb7tmn.

sphere that carves away the space outside a union of balls. The disadvantage of this type of surfaces is that they may not be tangent continuous.

References

[1] Computational Geometry Algorithms Library. `http://www.cgal.org`.

[2] D. Attali and J-D. Boissonnat. A linear bound on the complexity of the Delaunay triangulation of points on polyhedral surfaces. In *Proceedings of the seventh ACM symposium on Solid modeling and applications*, pages 139–146. ACM Press, 2002.

[3] C. Bajaj, H.Y. Lee, R. Merkert, and V. Pascucci. Nurbs based b-rep models for macromolecules and their properties. In *Proceedings of the fourth ACM symposium on Solid modeling and applications*, pages 217–228. ACM Press, 1997.

[4] C.L. Bajaj, V. Pascucci, A. Shamir, R.J. Holt, and A.N. Netravali. Dynamic maintenance and visualization of molecular surfaces. *Discrete Appl. Math.*, 127(1):23–51, 2003.

[5] J. Bloomenthal. Polygonization of implicit surfaces. *Comput. Aided Geom. Design*, 5(4):341–355, 1988.

[6] J-D. Boissonnat, D. Cohen-Steiner, and G. Vegter. Isotopic implicit surface meshing. In *ACM Symposium on Theory of Computing (STOC)*, pages 301–309, 2004.

[7] Jean-Daniel Boissonnat and Mariette Yvinec. *Algorithmic Geometry*. Cambridge University Press, UK, 1998. Translated by Hervé Brönnimann.

[8] H. Brönnimann, C. Burnikel, and S. Pion. Interval arithmetic yields efficient dynamic filters for computational geometry. *Discrete Applied Mathematics*, 109:25–47, 2001.

[9] H-L. Cheng, T.K. Dey, H. Edelsbrunner, and J. Sullivan. Dynamic skin triangulation. *Discrete Comput. Geom.*, 25:525–568, 2001.

[10] H.-L. Cheng and X.-W. Shi. Guaranteed quality triangulation of molecular skin surfaces. In *IEEE Visualization*, 2004.

[11] L. Chew. Guaranteed quality Delaunay meshing in 3d. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 391–393, 1997.

[12] S. Choi and N. Amenta. Delaunay triangulation programs on surface data. In *The 13th ACM-SIAM Symp. on Discrete Algorithms*, pages 135–136, 2002.

[13] M.L. Connolly. Analytical molecular surface calculation. *Journal of Applied Crystallography*, 16(5):548–558, Oct 1983.

[14] H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.*, 21:87–115, 1999.

[15] H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, 1990.

[16] H. Edelsbrunner and A. Üngör. Relaxed scheduling in dynamic skin triangulation. In *Japanese Conf. Comput. Geom.*, pages 135–151, 2002.

[17] Erich Hartmann. A marching method for the triangulation of surfaces. *The Visual Computer*, 14(3):95–108, 1998. ISSN 0178-2789.

[18] Leif Kobbelt. 3 -subdivision. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 103–112. ACM Press/Addison-Wesley Publishing Co., 2000.

[19] N.G.H. Kruithof and G. Vegter. Triangulating skin surfaces. Technical Report ECG-TR-244303-01, Rijksuniversiteit Groningen, 2003.

[20] Nico Kruithof and Gert Vegter. Approximation by skin surfaces. *Computer-Aided Design*, 36:1075–1088, 2004.

[21] W. Lorensen and H. Cline. Marching cubes: a high resolution 3d surface construction algorithm. *Comput. Graph.*, 21(4):163–170, 1987.

[22] D. Pedoe. *Geometry, a comprehensive course*. Dover Publications, New York, 1970.

[23] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *ACM SIGGRAPH symposium on Geometry processing*, pages 251–260, 2004.

[24] G.M. Treece, R.W. Prager, and A.H. Gee. Regularised marching tetrahedra: improved iso-surface extraction. *Computers and Graphics*, 23(4):583–598, 1999.