Dynamic Skin Triangulation *

(extended abstract)

Ho-Lun Cheng†, Tamal K. Dey‡, Herbert Edelsbrunner§ and John Sullivan¶

Abstract

This paper describes an algorithm for maintaining an approximating triangulation of a deforming surface in $\mathbb{R}^3$. The triangulation adapts dynamically to changing shape, curvature, and topology of the surface.

Keywords. Computational geometry, differential geometry, adaptive meshing, deformation, metamorphoses, algorithms, proofs.

1 Introduction

This paper develops a fully dynamic algorithm for maintaining a triangulation of a surface embedded in $\mathbb{R}^3$ that changes its local and global shape, curvature, and topology with time.

Motivation. Deforming surfaces arise in moving boundary problems of physical simulation, where they act as boundaries of spatial domains that grow and shrink with time. An example is the boundary between the solid and the liquid portions of metal during solidification [10]. Another is the phase boundary in a solid alloy that goes through the nucleation, growth and coarsening stages [1]. Moving boundaries also arise naturally in mold filling processes, both for metal and other materials [8]. Such physical processes are simulated through numerical computations facilitated by a mesh representing the boundary and/or domain. This mesh may be a 2-dimensional triangulation of the surface, or a 3-dimensional triangulation of space on one or both sides of the surface. The numerical methods require that the triangles and tetrahedra used in the triangulation be well-shaped, which usually means they have small aspect ratio, or equivalently they avoid small and large angles.

Skin Surfaces. The approach to deforming surfaces taken in this paper is based on the technical notion of skin surfaces, as introduced in [3]. The main reason for this choice is the existence of fast combinatorial algorithms based on the alpha shape theory [4]. A skin surface is defined by a finite collection of spheres in $\mathbb{R}^3$. We can think of the spheres as points with real weights, and we occasionally prefer this interpretation to avoid confusion with the various other types of spheres that arise in this paper. We derive an infinite family of spheres from the finite collection by convex combination and shrinking. The skin surface is the envelope of this family. Even though the family is infinite, the surface can be finitely described through a decomposition into a collection of quadratic surface patches. Each patch is the portion of a sphere or a hyperboloid lying inside a convex polyhedron. In each case, the sphere or hyperboloid and the containing polyhedron are defined by $k \leq 4$ weighted points (the original spheres). These polyhedra taken together form a finite tiling of space, which we refer to as the mixed complex. The correctness of the essentially combinatorial surface triangulation algorithm relies on the availability of exact geometric information, possibly in symbolic form. Most important in this context is the maximum curvature at a given surface point, which we show varies continuously over the surface and over space. This and other geometric information is readily computable from the mixed complex and the decomposition of the surface it defines.

Triangulation. The skin surface is represented by a 2-dimensional triangulation. We follow the convention in topology, where a triangulation means a simplicial complex whose underlying space is homeomorphic to the surface. The triangulation also approximates the surface. Specifically, its vertices lie on the surface and they are spread out depending on curvature. The algorithm maintains the triangulation through local restructuring operations:

*Research by all authors is partially supported by NSF under grant DMS-98-73945. Research by the first and third authors is also partially supported by ARO under grant DAAG55-98-1-0177. Research by the third author is also partially supported by NSF under grants CCR-96-19542 and CCR-97-12088. Research by the fourth author is also partially supported by NSF under grant DMS-97-27859.

†Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois.

‡Department of Computer Science, Ohio State University, Columbus, Ohio.

§Department of Computer Science, Duke University, Durham, and Raindrop Geomagic, Research Triangle Park, North Carolina.

¶Department of Mathematics, University of Illinois at Urbana-Champaign, Urbana, Illinois.

- it *moves vertices* in space to adapt the triangulation to changing shape,
- it *adds* and *removes vertices* to adapt the local density to the local maximum curvature,
- it *adjusts connectivity* provided by edges and triangles to reflect changing topology.

The local operations are automatic and follow the deformation of the surface dictated by the gradual change of the defining weighted points. The maximum curvature at each surface point is a single real number, so the adaptation of local density produces an isotropic triangulation. We gain flexibility by permitting the triangles to deviate somewhat from the equilateral shape. The deviation is measured as circumradius over length of the shortest edge, and the algorithm guarantees that this ratio never exceeds $\frac{Q^2}{2}$. Here Q is one of the constants on which the algorithm depends, the other being C. The two constants need to be chosen judiciously in order to guarantee the correctness of the algorithm. C controls how well the triangulation approximates the surface, and Q controls how far local density can deviate from strict inverse proportionality to local maximum curvature.

PART I. GEOMETRY

The three sections here introduce the skin surface, analyze its tangent and curvature behavior, and show that with a dense sampling we can triangulate the surface using the restricted Delaunay triangulation.

2 Skin Surfaces

The description of skin surfaces and their properties offered in this section is perhaps somewhat terse. The reader who wishes more background material is referred to [3] for the original introduction of skin surfaces, to [4] for a description of alpha shapes, and to [9] for a textbook in geometry that talks about the vector space of spheres used in the construction of skin surfaces.

Sphere algebra. Let $\hat{a} = (a, A)$ be the sphere with center $a \in \mathbb{R}^3$ and radius A. We require $A^2 \in \mathbb{R}$. For $A^2 < 0$ the radius is imaginary and we call $\hat{a}$ an imaginary sphere. Its *weighted (square) distance function* $\pi_{\hat{a}} : \mathbb{R}^3 \to \mathbb{R}$ is defined by $\pi_{\hat{a}}(x) = \|x - a\|^2 - A^2$; the original sphere is the zero-set of this function. We know how to add functions and how to multiply them by scalars. If we apply these operations to the $\pi_{\hat{a}}$ we get the vector space of functions of the form $\pi(x) = \gamma \cdot (\|x - p\|^2 - \beta)$, where $\beta, \gamma \in \mathbb{R}$ are scalars and $p \in \mathbb{R}^3$ is a point. The zero-set of π is the sphere with center p and radius $\sqrt{\beta}$.

We simplify notation by applying operations directly to spheres. In particular, we write $\hat{a} + \hat{b}$ for the zero-set of $\pi_{\hat{a}} + \pi_{\hat{b}}$ and $\gamma \cdot \hat{a}$ for the zero-set of $\gamma \cdot \pi_{\hat{a}}$. Using this notation, we can define what we mean by the *affine hull* and by the *convex hull* of a finite collection of spheres $\mathcal{A} = \{\hat{a}_1, \hat{a}_2, \ldots, \hat{a}_n\}$, namely

$$\begin{aligned} \operatorname{aff} \mathcal{A} &= \{\sum_{i=1}^{n} \gamma_i \cdot \hat{a}_i \mid \sum_{i=1}^{n} \gamma_i = 1\}, \\ \operatorname{conv} \mathcal{A} &= \{\sum_{i=1}^{n} \gamma_i \cdot \hat{a}_i \in \operatorname{aff} \mathcal{A} \mid \gamma_i \geq 0 \text{ for all } i\}. \end{aligned}$$

As an exercise, the reader may want to verify that if $\mathcal{A}$ contains only two spheres and they intersect in a common circle then the affine hull contains exactly all spheres passing through this circle. The convex hull contains the subset whose centers lie on the line segment connecting the centers of the two given spheres.

Besides adding and multiplying with a scalar, we need to be able to shrink spheres. For this purpose we define $\sqrt{\hat{a}} = (a, A/\sqrt{2})$, which is the zero-set of $2\pi_{\hat{a}} + A^2$. The application of the shrinking operation to all spheres in a family $\mathcal{F}$ is denoted as $\sqrt{\mathcal{F}} = \{\sqrt{\hat{a}} \mid \hat{a} \in \mathcal{F}\}$. The *skin* is the envelope of the spheres in the convex hull after shrinking,

$$\operatorname{skin} \mathcal{A} = \operatorname{env} \sqrt{\operatorname{conv} \mathcal{A}}.$$

It is also the boundary of the *body*, denoted as $\operatorname{body} \mathcal{A}$, which is the union of balls bounded by spheres in $\sqrt{\operatorname{conv} \mathcal{A}}$.

Mixed cells. The mixed cells are obtained from the corresponding weighted Voronoi and Delaunay faces. For a given finite collection of weighted points $\mathcal{A}$, the *Voronoi polyhedron* of $\hat{a} \in \mathcal{A}$ is the set of points x at least as close to $\hat{a}$ as to any other weighted point, $\nu_{\hat{a}} = \{x \in \mathbb{R}^3 \mid \pi_{\hat{a}}(x) \leq \pi_{\hat{b}}(x) \text{ for all } \hat{b} \in \mathcal{A}\}$. Two Voronoi polyhedra meet at most along a common piece of their boundary, and we define $\nu_{\mathcal{X}} = \bigcap_{\hat{a} \in \mathcal{X}} \nu_{\hat{a}}$ for every subset $\mathcal{X} \subseteq \mathcal{A}$. It is convenient to assume general position, in which case the dimension of each non-empty $\nu_{\mathcal{X}}$ is $\dim \nu_{\mathcal{X}} = 4 - \operatorname{card} \mathcal{X}$. In particular, $\nu_{\mathcal{X}}$ is a polyhedron, polygon, edge, vertex if $\operatorname{card} \mathcal{X} = 1$, 2, 3, 4, respectively. Each non-empty intersection of Voronoi polyhedra has a dual, which is geometrically realized as the convex hull of the (unweighted) points generating the polyhedra, $\delta_{\mathcal{X}} = \operatorname{conv} \{a \mid \hat{a} \in \mathcal{X}\}$. Assuming general position, the $\delta_{\mathcal{X}}$ are simplices, namely vertices, edges, triangles, tetrahedra. The collection of these simplices is referred to as the *Delaunay complex* of $\mathcal{A}$, although usually in the literature this term is reserved for the case of unweighted points.

Note that the dimensions of $\nu_{\mathcal{X}}$ and $\delta_{\mathcal{X}}$ are complementary, $\dim \nu_{\mathcal{X}} + \dim \delta_{\mathcal{X}} = 3$. Furthermore, they lie in orthogonal affine subspaces of $\mathbb{R}^3$. We use vector operations in $\mathbb{R}^3$ to construct the *mixed cell*,

$$\mu_{\mathcal{X}} = (\nu_{\mathcal{X}} + \delta_{\mathcal{X}})/2.$$

The dimension of $\mu_{\mathcal{X}}$ is always $\dim \mu_{\mathcal{X}} = \dim \nu_{\mathcal{X}} + \dim \delta_{\mathcal{X}} = 3$. Figure 1 shows examples of the four different types of mixed cells corresponding to different cardinalities of $\mathcal{X}$. The collection of mixed cells forms a face-to-

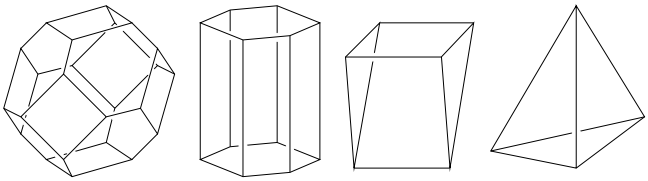

Figure 1: From left to right: a typical Voronoi polyhedron, a Voronoi polygon times a Delaunay edge, a Voronoi edge times a Delaunay triangle, a Delaunay tetrahedron.

face tiling of $\mathbb{R}^3$, which we call the *mixed complex*. Figure 4 shows the mixed complex defined by four points in the plane.

Skin patches. Within the mixed cell $\mu_{\mathcal{X}}$, the skin surface is completely determined by the at most four weighted points in $\mathcal{X}$ [3]. Specifically, it is the same as the envelope of the affine hull after shrinking all spheres, that is,

$$\operatorname{skin} \mathcal{A} \cap \mu_{\mathcal{X}} \quad = \quad \operatorname{env} \sqrt{\operatorname{aff} \mathcal{X}} \cap \mu_{\mathcal{X}}.$$

For $k = \operatorname{card} \mathcal{X} - 1 = 0$ or 3 the envelope of $\sqrt{\operatorname{aff} \mathcal{X}}$ is a sphere, and for $k = 1$ or 2 it is a hyperboloid of revolution. The hyperboloids have asymptotic double-cones with right opening angles. In each case, we define the *center* equal to the point $z_{\mathcal{X}}$ that is common to the affine subspace defined by $\nu_{\mathcal{X}}$ and that defined by $\delta_{\mathcal{X}}$. In the case of a hyperboloid this is the apex of the asymptotic double-cone, and in the case of a sphere it is the center. It may or may not belong to the mixed cell, and we have $z_{\mathcal{X}} \in \mu_{\mathcal{X}}$ iff $z_{\mathcal{X}} = \nu_{\mathcal{X}} \cap \delta_{\mathcal{X}}$.

If we translate the center to the origin and, in the case of a hyperboloid ($k = 1$ or 2), rotate so that the axis of symmetry is along the x_3-axis, we put the envelope into *standard form*. If R is the minimum distance from the origin to the envelope, then the equations of the envelopes for the sphere and the hyperboloids are

$$x_1^2 + x_2^2 + x_3^2 \quad = \quad R^2 \qquad (1)$$
$$x_1^2 + x_2^2 - x_3^2 \quad = \quad \pm R^2, \qquad (2)$$

We take the plus sign for the one-sheeted hyperboloid and the minus sign for the two-sheeted hyperboloid. The double-cone arises as the limiting case for $R = 0$. The three surfaces are illustrated in Figure 2.

Metamorphoses. A rather simple kind of deformation of the skin surface is generated by increasing the weight of every point in $\mathcal{A}$ in a uniform manner. We call this the *growth model* of deformation. It is generated by changing the original weight A^2 of the weighted point $\hat{a}$ to $A^2 + t$ at time t. Even though the mixed cells remain unaffected by the weight change, we observe all generic types of topological changes or metamorphoses that arise in general deformations. As indicated in Table 1, there are four types depending

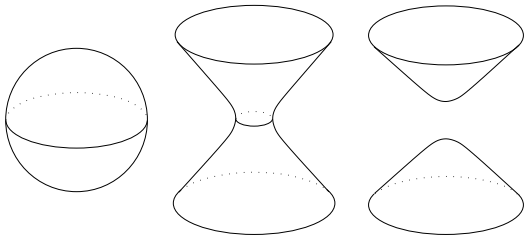

Figure 2: The sphere, the one-sheeted hyperboloid, and the two-sheeted hyperboloid.

k	type of metamorphosis/inverse
0	creating/annihilating a component
1	adding/removing a handle
2	closing/opening a tunnel
3	filling/starting a void

Table 1: The four types of generic metamorphoses that happen during growth/shrinking.

on $k = \operatorname{card} \mathcal{X} - 1$. By reversing time we get the inverse operations.

We can also reverse the orientation of the skin surface by finding another finite collection of weighted points that has the same skin and a complementary body. Specifically, there is a collection of spheres $\mathcal{B} = \mathcal{A}^{\perp}$ with $\operatorname{skin} \mathcal{A} = \operatorname{skin} \mathcal{B}$ and $\operatorname{body} \mathcal{A} \cup \operatorname{body} \mathcal{B} = \mathbb{R}^3$. Essentially, $\mathcal{B}$ contains a weighted point at every Voronoi vertex $b = \nu_{\mathcal{X}}$, and the weight is chosen as $B^2 = \|a - b\|^2 - A^2$ for every $\hat{a} \in \mathcal{X}$ [3]. When we revisit the metamorphoses listed in Table 1 and reinterpret them by what they do to the body of $\mathcal{B}$, we notice a symmetry between cases $k = 0, 3$ and another between cases $k = 1, 2$. In other words, there are only two basic types of metamorphoses. The first type is geometrically realized by a sphere appearing or disappearing. The limit configuration is a point, and in the growth model this is the center of the sphere. The second type of metamorphosis is geometrically realized by a two-sheeted hyperboloid flipping over to a one-sheeted hyperboloid, or vice versa. The limit configuration is a double-cone, and in the growth model this is the shared asymptotic double-cone of the two hyperboloids.

Time of change. An interesting question is when exactly do metamorphoses happen. We answer this in the context of the growth model of deformation by introducing certain subcomplexes of the Delaunay complex. In the literature, these subcomplexes are referred to as alpha complexes [4], but we use different notation and simply denote them by $K(t)$, where $t \in \mathbb{R}$ is time as before. Restrict each Voronoi polyhedron to within the generating sphere at time t, namely

$$\nu_{\hat{a}}(t) \quad = \quad \{x \in \nu_{\hat{a}} \mid \|x - a\|^2 \leq A^2 + t\}.$$

The complex $K(t)$ consists of all Delaunay simplices $\delta_{\mathcal{X}}$ for which the restricted Voronoi polyhedra have non-empty

intersection, that is, $\bigcap_{\hat{a}\in\mathcal{X}} \nu_{\hat{a}}(t) \neq \emptyset$. As t increases, $K = K(t)$ grows into a progressively larger subcomplex until eventually it is equal to the Delaunay complex. We sort the simplices in the order they enter the complex K. The resulting sequence of Delaunay simplices captures the evolution of the complex. Note that every prefix of the sequence is itself a complex. Whenever we have this property, we also have a fast algorithm for deciding how and when the homotopy type of K changes [2].

The underlying space of $K(t)$ and the body bounded by the skin at time t are homotopy equivalent [3]. It follows that the metamorphoses for the two structures happen at exactly the same moments in time, and these moments can be computed from the sequence of simplices. The case k corresponds to adding a k-dimensional simplex to K. However, not every addition of a new simplex corresponds to a metamorphosis. Only the addition of a Delaunay simplex $\delta_{\mathcal{X}}$ that contains its center $z_{\mathcal{X}}$ changes the homotopy type of K.

Sandwiching spheres. We close this section by stating a rather special property of skin surfaces heavily exploited in this paper. We mentioned already that the skin is the envelope of two families of spheres, one inside and the other outside the surface. As always we write $\mathcal{B} = \mathcal{A}^{\perp}$.

SANDWICH PROPERTY. For every point x on the skin of $\mathcal{A}$, there are unique spheres $S_x \in \sqrt{\operatorname{conv}\mathcal{A}}$ and $T_x \in \sqrt{\operatorname{conv}\mathcal{B}}$ that pass through x. Furthermore, S_x and T_x have the same radius.

We refer to S_x and T_x as the *sandwiching spheres* at x because they squeeze the surface flat in a neighborhood of x. They also limit the normal curvatures at x, and we will see in Section 3 that they in fact determine the maximum curvature. The fact that S_x and T_x are equally large follows from Lemma 7 in [3].

3 Continuity of Curvature

This section proves that the maximum curvature is continuous and observes a Lipschitz condition. We use this to control local density in the triangulation. This section also proves a Lipschitz condition for the normal direction.

Maximum curvature. For a point x and a tangent vector $\mathbf{t}_x$, the *normal curvature* is that of a geodesic passing through x in the direction $\mathbf{t}_x$. The *maximum curvature* is the function $\kappa : F \to \mathbb{R}$ that maps $x \in F$ to the maximum normal curvature at x. For a hyperboloid, the minimum curvature is measured alongside within planes containing the symmetry axis, and the maximum curvature is measured in the orthogonal direction. Explicit expressions for κ are easy to compute [6, Chap. 14]. For the sphere and the hyperboloids in standard form (1) and (2), the maximum curvatures are $\kappa = 1/R$ and $\kappa = 1/\sqrt{\pm R^2 + 2x_3^2}$, where we take the plus sign for one-sheeted hyperboloids and the minus sign for two-sheeted hyperboloids. By plugging $\pm R^2 = x_1^2 + x_2^2 - x_3^2$ into the equation for hyperboloids we see that the maximum curvature at x is one over the distance of x from the origin. This implies that points with constant maximum curvature form spherical shells around the origin.

ISO-CURVATURE LEMMA. Every point $x \in \mathbb{R}^3$ belongs to exactly one hyperboloid in standard form, and its maximum curvature in that hyperboloid is $\kappa(x) = \frac{1}{\|x\|}$.

For either type of hyperboloid, the curvature is everywhere at most $\frac{1}{R}$. For the one-sheeted hyperboloid, R is also the radius of the smallest circle around the neck of the hour-glass. For the two-sheeted hyperboloid, R is also half the smallest distance between the two sheets.

Curvature variation. To prove that κ varies continuously over the skin surface, we consider the two infinite families of spheres that define the skin as their common envelope. For a finite set of spheres $\mathcal{A}$, let $S = \sqrt{\operatorname{conv}\mathcal{A}}$ and $T = \sqrt{\operatorname{conv}\mathcal{A}^{\perp}}$. The skin of $\mathcal{A}$ is $F = \operatorname{env} S = \operatorname{env} T$. The family S defines F from the inside, and T defines it from the outside. For a point $x \in F$, there are unique spheres $S_x \in S$ and $T_x \in T$ that pass through x. We make essential use of the Sandwich Property.

CURVATURE SANDWICH LEMMA. For every point $x \in F$, the reciprocal of the maximum curvature, $\frac{1}{\kappa(x)}$, is the common radius of the sandwiching spheres S_x and T_x.

The sandwiching spheres vary continuously with x. The Curvature Sandwich Lemma thus implies that the maximum curvature varies continuously over the skin surface. We strengthen the result by proving that κ varies only slowly with the location. It is convenient to define $\varrho(x) = 1/\kappa(x)$, and for reasons that will become clear later we refer to $\varrho : F \to \mathbb{R}$ as the *length scale*. Specifically, we prove that the difference in ϱ is at most the Euclidean distance in $\mathbb{R}^3$.

CURVATURE VARIATION LEMMA. For all points x and y on the skin surface we have $|\varrho(x) - \varrho(y)| \leq \|x - y\|$.

The proof does not require x and y to belong to the same skin surface. The inequality also holds if x and y belong to different skin surfaces in the one-parameter growth model of deformation.

Normal variation. The tangent or C^1-continuity of the skin surface follows from the Sandwich Property. We strengthen this result by proving a Lipschitz condition for the normal vectors. Specifically, we prove an upper bound that relates the angle between two normal vectors at points x, y to the Euclidean distance between and the length scale at the points. The outward directed unit normal vector at $x \in F$ is denoted as $\mathbf{n}_x$, and the angle between two normals is $\angle\mathbf{n}_x\mathbf{n}_y = \arccos(\mathbf{n}_x \cdot \mathbf{n}_y)$. In proving the upper bound,

we consider again the one-parameter family of skin surfaces generated by increasing square radii with time. For points $x = (x_1, x_2, x_3)$ on a sphere in standard form the unit normals are $\mathbf{n}_x = \pm x/\|x\|$, and for points x on a hyperboloid in standard form they are $\mathbf{n}_x = \pm(x_1, x_2, -x_3)/\|x\|$. In both cases, the normals are the same along a line passing through the origin, and they vary with the angle as we rotate the point about the origin.

NORMAL VARIATION LEMMA. Let $x, y \in F$ with $\varrho(y) \leq \varrho(x)$ and $\|x - y\| < \varrho(x)$. The angle between the surface normals at x and y is $\angle \mathbf{n}_x \mathbf{n}_y \leq \arcsin \frac{\|x-y\|}{\varrho(x)}$.

Similar to the proof of the Curvature Variation Lemma, the proof of the Normal Variation Lemma does not require that x and y belong to the same skin surface.

4 Triangulation

A finite set $V \subseteq F$ is an *ε-sampling* if for every point $x \in F$ there is a vertex $a \in V$ whose distance from x is $\|x - a\| < \varepsilon \cdot \varrho(x)$. The goal of this section is to prove that an ε-sampling for F defines a restricted Delaunay triangulation homeomorphic to the skin surface, provided

(I) $0 \leq \varepsilon \leq \varepsilon_0$, where $\varepsilon_0 = 0.279\ldots$ is a root of $f(\varepsilon) = 2\cos(\arcsin \frac{2\varepsilon}{1-\varepsilon} + \arcsin \varepsilon) - \frac{2\varepsilon}{1-\varepsilon}$.

Restricted Delaunay triangulation. Let V be a finite set of points on the skin surface. We refer to these points as vertices. Since $V \subseteq \mathbb{R}^3$, the Voronoi polyhedron ν_a is defined for each $a \in V$. The corresponding *restricted Voronoi polygon* is the intersection with the skin surface, $F \cap \nu_a$. The *restricted Delaunay triangulation* is the nerve of the collection of restricted polygons:

$$D_V \quad = \quad \{\operatorname{conv} U \mid U \subseteq V, F \cap \bigcap_{a \in U} \nu_a \neq \emptyset\}.$$

We assume general position and in particular that there are no four restricted Voronoi polygons with non-empty common intersection. It follows that $D = D_V$ is a collection of vertices, edges, triangles but contains no tetrahedra. By construction, D is a simplicial complex. The goal of this section is to prove that, for ε satisfying Condition (I), D is a triangulation of F. Following the standard topology terminology, this means the underlying space of D is homeomorphic to F. As shown in [5], it suffices to prove that every non-empty common intersection of restricted Voronoi polygons is a closed ball of the appropriate dimension, namely 3 minus the number of polygons. If this is the case we say D has the *closed ball property*.

Summary. We omit the somewhat involved proof that for ε satisfying Condition (I), the restricted Voronoi diagram of an ε-sampling V has the closed ball property. Assuming this property, the result of [5] implies that the underlying space of the restricted Delaunay triangulation is homeomorphic to the skin surface.

GENERAL HOMEOMORPHISM THEOREM. The restricted Delaunay triangulation of an ε-sampling triangulates the skin surface.

For the purpose of changing the topology of the skin surface we will rely on point distributions that locally violate the ε-sampling condition. We will give a separate proof of the closed ball property in Section 10 and thus obtain a Special Homeomorphism Theorem for such distributions.

PART II. ALGORITHM

The algorithm maintains the triangulation of a deforming skin surface dynamically by adapting geometric location to shape, density to curvature, and connectivity to topology. It can be used to construct a triangulation by starting with the empty triangulation and growing components from nothing.

5 Shape Adaptation

Let $t_0 < t_1$ be moments in time and D_0, D_1 the corresponding restricted Delaunay triangulations. The algorithm updates D_0 locally and changes it to D_1. This section describes the overall algorithm and presents the details for adapting the triangulation to the changing shape of the surface.

Moving vertices. The intuition for moving vertices is taken from Morse theory, which considers differential structures that arise in sweeping out a smooth manifold [7]. The skin surface is the cross-section at a moment in time during the sweep, and the manifold is the stack of cross-sections in the time direction. In other words, the manifold is the graph of $M : \mathbb{R}^3 \to \mathbb{R}$ that maps a point x to the time t at which x belongs to the surface $F(t)$. Hence $F(t) = M^{-1}(t)$. A metamorphosis of F corresponds to a critical point of M. For cross-sections in a time interval $[t_0, t_1]$ that is free of critical points, we can construct a 1-parameter family of diffeomorphisms from the integral lines of the ordinary differential equation defined by the gradient $\operatorname{grad} M : \mathbb{R}^3 \to \mathbb{R}^3$ that maps a point x to the vector $\operatorname{grad} M(x) = \frac{\partial M}{\partial t}(x)$. These diffeomorphisms $\varphi_i : F(t_0) \to F(t_i)$, with $t_i \in [t_0, t_1]$, can be composed to diffeomorphically connect any two cross-sections in the time interval,

$$\varphi_{ij} = \varphi_j \circ \varphi_i^{-1} : F(t_i) \to F(t_j).$$

The step from time t_i to time t_j thus amounts to moving each vertex $a \in D_i$ along its integral line to $a' = \varphi_{ij}(a) \in D_j$. In the growth model of deformation, the integral lines are pieces of straight lines and hyperbolas, as illustrated in Figure 3. To see this note that $(2x_1, -2x_3)$ are the normal vectors of the

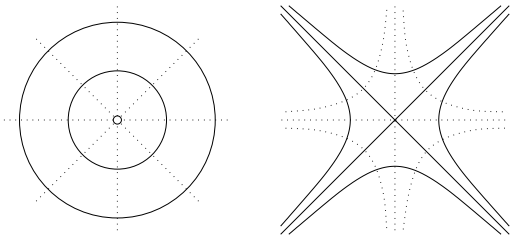

Figure 3: Dotted integral lines of a solid growing circle and a solid growing hyperbola.

family of hyperbolas $x_1^2 - x_3^2 = \pm R^2$, and that $(2x_3, 2x_1)$ are the normal vectors of the family $x_1 x_3 = \pm R^2/2$ obtained by rotating the first family through an angle of $\frac{\pi}{4}$. The 3-dimensional picture is obtained by revolving the hyperbolas in Figure 3 about the x_3-axis. The first family of hyperbolas turns into the 1-parameter family of hyperboloids described by equation (2). The second family turns into a 2-parameter family of hyperbolas each orthogonal to every hyperboloid in the 1-parameter family.

Algorithmic time-warp. Vertices move continuously along their integral lines, but updating them continuously is computationally infeasible. The common escape from this dilemma is the time-slicing method, which takes discrete time steps and advances all vertices from time t_0 to time t_1 without intermediate stop. There are drawbacks to time-slicing related to the difficulty of choosing the right step size. We thus follow an alternative approach and take different time steps at different locations. This is done by prioritizing the operations that occur at discrete moments in time, which are edge flips, edge contractions, vertex insertions, and metamorphoses. Coordinate updates are done lazily, moving a vertex when and only when it is used by one of the other four operations. This results in a time-warped surface with different pieces reflecting the state at different times. To bring the entire surface to the present time, we simply update all the vertex coordinates, and by assumed correctness of the prioritization this requires no other changes in the triangulation.

A any moment in time t, we consider the collection of possible next operations. Let $t_i > t$ be the time such an operation τ_i would happen if the vertices moved along integral lines and no other operations preceded τ_i. We store the τ_i in a priority queue ordered by time.

`loop` $\tau_i =$ NEXTOP; $D =$ APPLY(τ_i) `forever`.

Function APPLY changes D according to τ_i, and simultaneously updates the priority queue by inserting new operations made possible by the changes caused by τ_i. The changes may make some of the operations in the priority queue inapplicable. We use a lazy strategy that checks an operation when it reaches the top of the priority queue.

`Operation` NEXTOP:
  `repeat` $\tau =$ EXTRACTMIN `until` ISOK(τ);
  `return` τ.

Determining when exactly an operation τ_i matures in the future is computationally fairly expensive, and so is the correct ordering of operations in time. We plan to discuss approximate ordering methods that alleviate the cost in a later paper.

6 Curvature Adaptation

This section focuses on the density adaptation algorithm. The method is straightforward, but we need some geometric analysis to convince ourselves that it is correct.

Invariants. The goal of the algorithm is to locally triangulate with edges and triangles of size roughly proportional to the length scale, which Section 3 defined as one over the maximum curvature, $\varrho(x) = 1/\kappa(x)$. The *size* of an edge ab is defined to be half its length, $R_{ab} = \|a - b\|/2$. The *size* of a triangle abc is the radius of the circumcircle and denoted as R_{abc}. Since edges and triangles are not points, it is ambiguous which length scale exactly they should follow. For edges we worry about them getting too short, so we compare size with the maximum length scale, and for triangles we worry about them growing too large, so we compare size with the minimum length scale:

$$\begin{aligned} \varrho_{ab} &= \max\{\varrho(a), \varrho(b)\}, \\ \varrho_{abc} &= \min\{\varrho(a), \varrho(b), \varrho(c)\}. \end{aligned}$$

The algorithm is formulated using two positive constants, C and Q. Roughly, C controls how closely the triangulation approximates the skin surface, and Q controls the quality of the triangles. The following two inequalities are maintained as invariants, which we refer to as the Lower and the Upper Size Bounds:

[L] $\frac{R_{ab}}{\varrho_{ab}} > \frac{C}{Q}$ for every edge $ab \in D$.

[U] $\frac{R_{abc}}{\varrho_{abc}} < CQ$ for every triangle $abc \in D$.

It is not necessary to explicitly check for long edges and small triangles. This is because an edge of size $R_{ab} \geq CQ \cdot \varrho_{ab}$ belongs to two triangles that both violate [U]. Symmetrically, a triangle of size $R_{abc} \leq \frac{C}{Q} \cdot \varrho_{abc}$ has at least two edges that violate [L]. Appropriate values of C, Q will be determined in the analysis of the algorithm but we can already anticipate $C = 0.1$, $Q = 1.65$ as a feasible assignment.

Minimum angle. The smallest angle is a measure of triangle quality. It achieves its maximum for the equilateral triangle, for which it is $\frac{\pi}{3}$. Triangles that satisfy both Size Bounds cannot have arbitrarily small angles.

MINIMUM ANGLE LEMMA. A triangle that satisfies [L] and [U] has minimum angle larger than $\arcsin \frac{1}{Q^2}$.

The Minimum Angle Lemma suggests that we choose Q as small as possible, contingent upon satisfying all constraints needed to prove the algorithm correct. For $Q = 1.65$ the minimum angle is larger than $21.54\ldots^\circ$, and the maximum angle is smaller than $136.90\ldots^\circ$.

Enforcement. The algorithm enforces the two invariants by contracting short edges and inserting vertices near the barycenters of large triangles. Let abc be a triangle that gets too large, that is, $R_{abc} = CQ \cdot \varrho_{abc}$ at time t_i. To remedy the violation of the Upper Size Bound, we add the restricted Voronoi vertex x dual to abc as a new vertex to the triangulation. A vertex insertion may cause new violations of the Upper Size Bound and thus trigger additional vertex insertions. We thus apply them in a loop until no offending triangles remain:

```
void VertexInsertion:
  while ∃ triangle abc violating [U] do
    Add(x, abc)
  endwhile.
```

Consider next an edge ab that gets too short, that is, $R_{ab} = \frac{C}{Q} \cdot \varrho_{ab}$ at time t_j. To remedy the violation of the Lower Size Bound, we contract ab by removing b from the triangulation. The removal of b may possibly create new edges violating [L], and it can certainly create triangles violating [U]. We repair the triangulation in two nested loops.

```
void EdgeContraction:
  while ∃ edge ab violating [L] do
    Remove(b); VertexInsertion
  endwhile.
```

7 Topology Adaptation

The way the skin surface is connected can change during deformation. This section studies when, where, and how these changes happen. It also describes how we locally modify the general sampling strategy to avoid the computational impossibility of sampling infinitely many points accumulating at locations of infinite curvature.

Growth model of deformation. We recall the growth model of deformation defined by changing the square radius of a sphere (a, A) from A^2 at time 0 to $A^2 + t$ at time $t \in \mathbb{R}$. Computationally, this is the simplest kind of deformation because it keeps the mixed complex invariant. Each mixed cell contains a possibly empty sphere or hyperboloid patch of the skin surface. After normalization, the equation of the sphere or hyperboloid at time t is

$$x_1^2 + x_2^2 \pm x_3^2 \quad = \quad \pm R^2 + t. \tag{3}$$

Compare this with equations (1) and (2). A metamorphosis happens when the right-hand side vanishes at time $t = \mp R^2$, and it happens at the center but only if the center lies in the interior of its mixed cell. If the center lies outside, the portion of the sphere or hyperboloid that passes through the center is not part of and thus does not affect the skin surface. The special case where the center lies on the boundary of its mixed cell is interesting. In this case the metamorphosis does not happen, but we still have to modify the sampling strategy because curvature grows beyond every bound.

Using local considerations, we can reduce the list of metamorphoses to the four given in Table 1. Cases $k = 0, 3$ correspond to an appearing/disappearing sphere. Cases $k = 1, 2$ correspond to switching a hyperboloid from two sheets to one, or vice versa. In each case, we can interpret the center as a critical point of the map $M : \mathbb{R}^3 \to \mathbb{R}$ whose level sets $M^{-1}(t)$ are the skin surfaces at time t. Cases $k = 0, 3$ correspond to minima and maxima, and Cases $k = 1, 2$ to two types of saddle points. The gradient of M vanishes at all these points and also at centers that lie on the boundary of their mixed cells. The latter centers correspond to degenerate critical points in the sense that an arbitrarily small perturbation of M suffices to turn them into regular points.

Hot spots. Common to every metamorphosis is the local drop in length scale, which reaches zero at the moment and point of the metamorphosis. Let H be a positive real number. The *hot portion* of the skin surface F is the set of points with length scale H or smaller,

$$F_H \quad = \quad \{x \in F \mid \varrho(x) \leq H\}.$$

By the Iso-curvature Lemma, we have $\varrho(x) \leq H$ only if x is sufficiently close to the center of a sphere or hyperboloid. Let $z_\mathcal{X}$ be such a center. We call the ball $\beta_\mathcal{X} = \{y \in \mathbb{R}^3 \mid \|y - z_\mathcal{X}\| \leq H\}$ the *hot ball* of $\mathcal{X}$. It is relevant only inside its mixed cell. The union of hot balls, each clipped to within its mixed cell, is the *hot portion* of space, denoted as $\mathbb{R}^3_H$.

HOT SPOT LEMMA. $F_H = F \cap \mathbb{R}^3_H$.

Depending on whether centers lie inside or outside their mixed cells, the hot portion of space is locally a union or intersection of hot balls. The mixed complex decomposes this union and intersection into convex pieces, as illustrated in Figure 4. The common radius of all hot balls is H. As long as none of the centers lies on the boundary of its mixed cell, we can eliminate any overlap by decreasing H while keeping it positive. We will shortly discover that an even stronger separation property between hot balls is needed to prevent edges of the triangulation reaching from one hot sphere to another, which can be achieved e.g. by choosing H equal to half the value that guarantees pairwise disjointness. A center on a mixed cell boundary has probability zero and is considered a degenerate case. We simplify the discussion by assuming the non-degenerate case.

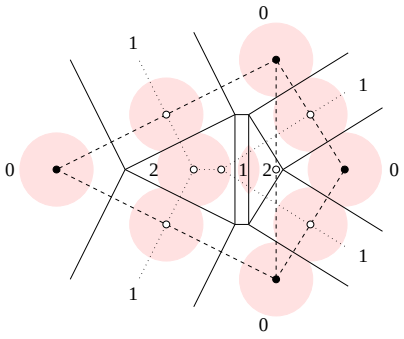

Figure 4: Dotted Voronoi diagram, dashed Delaunay triangulation, solid mixed complex, hollow centers, and shaded hot portion of space.

Time for change. The hot portion is more difficult to triangulate than the rest of the skin surface. One reason is the metamorphosis, another is the accumulation of vertices in a small region. The sphere case is relatively harmless, because the area decreases at the same rate as the density requirement increases. Indeed, a constant number of vertices suffices to shrink a sphere to arbitrarily small size. The case of a hyperboloid that approaches its limiting double-cone is more problematic, because the number of vertices near the two tips grows beyond every bound. To circumvent the computational impossibility of sampling infinitely many points, we change the sampling strategy on and inside the hot balls. We give up on ε-sampling to get a sparse sampling, but we preserve the closed ball property. The triangulation algorithm remains oblivious of the changed sampling density and keeps constructing the restricted Delaunay triangulation.

Consider a 2-sheeted hyperboloid and translate time such that the metamorphosis happens at time $t = 0$. The hyperboloid enters its hot ball at time $-H^2$, turns into a double-cone at time 0, and leaves the hot ball as a 1-sheeted hyperboloid at time H^2. The special sampling strategy that

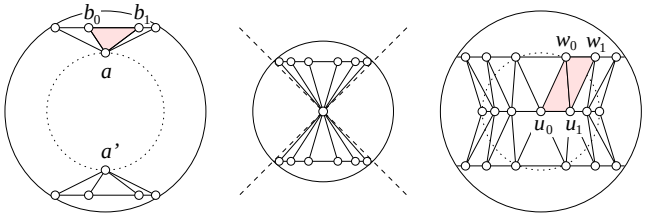

Figure 5: Head-on view of start, middle, end configurations generated by special sampling taking a 2-sheeted to a 1-sheeted hyperboloid. The hot sphere is solid and the sphere that triggers the metamorphosis is dotted.

allows us to go through this motion depends on a parameter $0 < h < 1$. Special sampling begins at time $t_0 = -H^2h^2$ when the 2-sheeted hyperboloid enters the ball of radius hH, and it ends at time $t_1 = H^2h^2$ when the 1-sheeted hyperboloid leaves that ball, as shown in Figure 5. At time t_0, the hyperboloid intersects the hot sphere in two *hot circles*. The shape adaptation algorithm moves these circles along their integral lines, which implies that they grow from radius $R_0 = H\sqrt{(1-h^2)/2}$ at time t_0 to radius $R_1 = H\sqrt{(1+h^2)/2}$ at time t_1. Simultaneously, the distance between the two circles decreases from $2R_1$ to $2R_0$. We define the *hot sphere* so it passes through the two hot circles. At times t_0 and t_1, it is the boundary of the hot ball, but in the open time interval between t_0 and t_1 it is cocentric and smaller than that boundary. General sampling applies outside the hot sphere and special sampling applies on and inside that sphere.

PART III. ANALYSIS

The next three sections analyze the algorithm and the triangulations it creates. Section 8 studies questions related to sampling density, Section 9 focuses on the scheduling aspects of the algorithm, and Section 10 examines the topology adaptation algorithm.

8 Sampling Density

Conditions for the constants C and Q are derived in order to prove the curvature adaptation algorithm in Part II is correct.

Conditions. We prove that point insertions do not generate edges that violate the Lower Size Bound. That proof requires that Q not be too large. We also prove that the vertices indeed form an ε-sampling, with ε satisfying Condition (I). The closed ball property established in Section 4 then implies that the produced triangulation is homeomorphic to the skin surface. That proof relies on the quality of the approximation, which is guaranteed by the algorithm provided CQ is not too large. For ease of reference we collect the conditions before deriving them in the analysis.

(II) $Q^2 - 4CQ - 2 > 0.$

(III) $\frac{\delta^2}{(1+\delta)^2} - \frac{\delta^4}{4} > C^2Q^2,$

where $\delta = \varepsilon - \frac{C(\varepsilon+1)}{Q}$. We get (II) and (III) as sufficient conditions for the proofs of the No-Short-Edge Lemma and the Sampling Lemma, which are omitted in this extended abstract. Condition (II) is equivalent to $Q > 2C + \sqrt{4C^2+2}$. Assuming $\varepsilon = \varepsilon_0 = 0.279\ldots$, we can satisfy Conditions (II) and (III) by setting $C = 0.1$ and $Q = 1.65$. In this case $\delta = 0.201\ldots$. Small improvements are possible.

Short edges. An edge contraction may perhaps cause other edge contractions, but this cannot go on forever because we eventually violate the Upper Size Bound. Similarly, a vertex insertion may cause other vertex insertions, but this cannot

go on forever because we eventually violate the Lower Size Bound. It is possible that an edge contraction causes vertex intersections, but a vertex insertion cannot cause edge contractions. This is because a vertex insertion cannot decrease the size of an edge below the allowed threshold, and it prevents infinite loops in spite of the algorithm's partially conflicting efforts to simultaneously avoid short edges and large triangles. Let abc be the triangle that causes the addition of the dual restricted Voronoi vertex $x \in F$.

NO-SHORT-EDGE LEMMA. Every edge xy created during the addition of x has size larger than $\frac{C}{Q} \cdot \varrho_{xy}$.

Maintaining density. We show that the algorithm for curvature adaptation maintains the ε-sampling property of the vertex set. Recall that this means that for every point $x \in F$ there is a vertex $p \in V$ whose distance from x is $\|p - x\| < \varepsilon \cdot \varrho(x)$. The constant ε is to be chosen so it satisfies Condition (I).

It is interesting to see that the two Size Bounds by themselves are too weak to imply ε-sampling. We can put four points near a pole of a sphere in such a way that all four triangles and six edges satisfy [L] and [U]. Nevertheless, the boundary of the tetrahedron is a miserably inadequate approximation of the sphere surface. We argue that the algorithm cannot get to this problematic state, because on the way it would have to temporarily violate at least one of the two Size Bounds. In other words, we use continuity in time to prove the claim on sampling. In stating the result, we assume the skin surface deforms continuously with time. Let $t_0 < t_1$ be two points in time. We write $F(t)$ for the skin surface at time t and V_0, V_1 for the vertex sets at times t_0, t_1.

SAMPLING LEMMA. If V_0 is an ε-sampling of $F(t_0)$ then V_1 is an ε-sampling of $F(t_1)$.

9 Scheduling

The overall algorithm deforms the skin surface by executing operations ordered in time. Some of these operations require others to repair the damage, and these are executed following a partial rather than a total order. As a general rule, total ordering is more expensive but easier to prove correct than partial ordering. This section reviews all operations and discusses their treatment by the scheduling algorithm.

Total and partial ordering. Operations triggered by the motion of the skin surface are ordered in time. We have five types, namely coordinate updates, edge flips, edge contractions, vertex insertions, metamorphoses. Vertex coordinates change continuously with time, and we avoid most of the related computational expense by updating them when and only when they are used by other operations. The last four operations are discrete events that are stored in a priority queue ordered by time. The moment in time when an edge flip, edge contraction, vertex insertion matures is a root of a continuous function. In the growth model of deformation, the moment in time when a metamorphosis matures is predictable from the ordering of Delaunay simplices described in Section 2. For more general deformations, the time of a metamorphosis also becomes more difficult to compute as a root of a continuous function.

Each operation other than the coordinate update and the edge flip is further decomposed into a sequence of operations. Conceptually, such a sequence is executed at an instance, while time stands still. We can therefore not resort to time for a global ordering mechanism. The operations in each sequence are therefore scheduled following a partial rather than a total order. The most frequently executed operation is the edge flip. The choice of constants C and Q guarantees that the restricted Voronoi diagram has the closed ball property at all times, even in the middle of an edge contraction. We would therefore expect that a simple iteration of edge flips will suffice to restore the restricted Delaunay triangulation. While this is easy to prove for point additions, it is possibly incorrect for point removals. This is why we resort to the more complicated edge flipping algorithm akin to deleting a point from a two-dimensional Delaunay triangulation, which is described in the full paper.

10 Metamorphoses

This section analyzes the point configurations generated by special sampling. Recall that $Hh < H$ is the length scale threshold that triggers the start and end of special sampling. In the forward direction we start with a 2-sheeted hyperboloid that enters the ball with radius Hh around its center, and we end with a 1-sheeted hyperboloid that exits the same ball. In the backwards direction the events are the same in reverse order.

Sizes at transition. Refer to the double-cup shown in Figure 5. The $\ell + 1$ points on one sheet form a regular ℓ-sided cup. The ℓ vertices of the base lie on the hot circle with radius $R_0 = H\sqrt{(1-h^2)/2}$, which lies in a plane at distance $R_1 = H\sqrt{(1+h^2)/2}$ from the center. Note that $R_0^2 + R_1^2 = H^2$. Define $b = b_i$ and $c = b_{i+1}$, with indices modulo ℓ. Independent of the index, the size are

$$\begin{aligned} R_{ab} &= \frac{1}{2} \cdot \sqrt{2R_1(R_1 - Hh)}, \\ R_{bc} &= R_0 \cdot \sin\frac{\pi}{\ell}, \\ R_{abc} &= \frac{R_1(R_1 - Hh)}{\sqrt{2R_1(R_1 - Hh) - R_0^2 \sin^2\frac{\pi}{\ell}}}. \end{aligned}$$

Next refer to the cylinder-with-a-waist shown in Figure 5. The $3m$ points form three parallel regular m-gons. The distance between two contiguous planes is R_0, and the circumradii of the three m-gons are R_1, Hh, R_1. Define $u =$

$u_i, v = u_{i+1}, w = w_i, x = w_{i+1}$, with indices modulo m. Independent of the index, the sizes are

$$\begin{aligned}
R_{uv} &= Hh \cdot \sin\frac{\pi}{m},\\
R_{wx} &= R_1 \cdot \sin\frac{\pi}{m},\\
R_{vw} &= \frac{1}{2} \cdot \sqrt{2R_1(R_1 - Hh \cdot \cos\frac{\pi}{m})},\\
R_{uvw} &= \frac{R_1(R_1 - Hh \cdot \cos\frac{\pi}{m})}{\sqrt{H^2 - 2R_1Hh \cdot \cos\frac{\pi}{m} + H^2h^2 \cdot \cos^2\frac{\pi}{m}}},\\
R_{vwx} &= \frac{R_1(R_1 - Hh \cdot \cos\frac{\pi}{m})}{\sqrt{R_1^2 - 2R_1Hh \cdot \cos\frac{\pi}{m} + R_1^2 \cdot \cos^2\frac{\pi}{m}}}.
\end{aligned}$$

Smooth transition. We derive necessary and sufficient conditions for h, ℓ, m that guarantee a smooth transition from the general to the special sampling strategy. By this we mean that the configurations at the start of a metamorphosis is an ε-sampling and satisfies both Size Bounds. At the end of the metamorphism, the Size Bounds are enforced by eliminating offending edges and triangles through edge contraction and vertex insertion. The result is a triangulation whose vertex set is an ε-sampling of the surface.

The length scale at the vertices a, u, v is Hh, and that at b, c, w, x is H. The Lower and Upper Size Bounds are therefore equivalent to $R_{ab}, R_{bc}, R_{uv}/h, R_{wx}, R_{vw} > \frac{C}{Q} \cdot H$ and $R_{abc}, R_{uvw}, R_{vwx} < CQ \cdot Hh$. The inequalities for R_{vw}, R_{uv}, R_{uvw} are redundant because $R_{ab} < R_{vw}$, $R_{wx} < R_{uv}/h$, $R_{uvw} < R_{vwx}$ for all $h < 1$. In addition to requiring that the triangles satisfy the Upper Size Bound, it is convenient to also require that their radii are less than twice the locally allowed edge length. This extra requirement implies that after adding points on and inside the hot sphere, all remaining old vertices lie outside that sphere. We thus have the following two conditions:

(IV) $\frac{R_{ab}}{H}, \frac{R_{bc}}{H}, \frac{R_{wx}}{H} > \frac{C}{Q}$,

(V) $\frac{R_{abc}}{H}, \frac{R_{vwx}}{H} < \min\{Q, \frac{2}{Q}\} \cdot Ch$.

Conditions (I) to (V) are satisfied for $\varepsilon = 0.279, C = 0.1, Q = 1.65, h = 0.97, \ell = 5, m = 30$. We summarize the results assuming this assignment of constants.

TRANSITION LEMMA. The triangulation at the start of a metamorphosis satisfies the two Size Bounds and its vertex set is an ε-sampling of the skin surface.

As mentioned earlier, the same does not automatically hold for the end configurations of metamorphoses, but it can be enforced algorithmically. The purpose of bounding the size of triangles in Condition (V) by $\frac{2}{Q} \cdot Ch$ is to guarantee that the algorithm given in Section 7 constructs the special configurations without having to search for remaining old vertices inside the hot sphere. To prove this algorithm correct, we also need to show that these configurations are part of the restricted Delaunay triangulation. Proofs are omitted.

Summary. The full version of this paper establishes that the closed ball property of the restricted Voronoi diagram is maintained even inside the hot spheres that guide the algorithm through the various metamorphoses.

SPECIAL HOMEOMORPHISM THEOREM. The restricted Delaunay triangulation of the points chosen by special sampling triangulates the skin surface inside each hot sphere.

Together with the General Homeomorphism Theorem this implies that we have a triangulation of the skin surface at all times.

References

[1] H. I. AARONSON, ED. *Lectures on the Theory of Phase Transformation.* American Institute of Mining, Metallurgical and Petroleum Engineer, New York, 1975.

[2] C. J. A. DELFINADO AND H. EDELSBRUNNER. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Comput. Aided Geom. Design* **12** (1995), 771–784.

[3] H. EDELSBRUNNER. Deformable smooth surface design. *Discrete Comput. Geom.* **21** (1999), 87–115.

[4] H. EDELSBRUNNER AND E. P. MÜCKE. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.

[5] H. EDELSBRUNNER AND N. R. SHAH. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.* **7** (1997), 365–378.

[6] A. GRAY. *Modern Differential Geometry of Curves and Surfaces.* CRC Press, Boca Raton, Florida, 1993.

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

[8] T. E. MORTHLAND, P. E. BYRNE, D. A. TORTORELLI AND J. A. DANTZIG. Optimal riser design for metal castings. *Metallurgical and Material Transactions* **26B** (1995), 871–885.

[9] D. PEDOE. *Geometry: a Comprehensive Course.* Dover, New York, 1988.

[10] N. PROVATAS, N. GOLDENFELD AND J. A. DANTZIG. Adaptive grid methods in solidification microstructure modeling. *Phys. Rev. Lett.* **80** (1998), 3308–3311.