Dynamic Skin Triangulation*

H.-L. Cheng,¹ T. K. Dey,² H. Edelsbrunner,³ and J. Sullivan⁴

¹Department of Computer Science, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA

²Department of Computer Science, Ohio State University,
Columbus, OH 43210, USA
tamaldey@cis.ohio.state.edu

³Department of Computer Science, Duke University,
Durham, NC 27708, USA
and
Raindrop Geomagic,
Research Triangle Park,
NC 27709, USA

⁴Department of Mathematics, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA

Abstract. This paper describes an algorithm for maintaining an approximating triangulation of a deforming surface in $\mathbb{R}^3$. The surface is the envelope of an infinite family of spheres defined and controlled by a finite collection of weighted points. The triangulation adapts dynamically to changing shape, curvature, and topology of the surface.

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.

* Research by all authors was partially supported by NSF under Grant DMS-98-73945. Research by the first and third authors was also partially supported by ARO under Grant DAAG55-98-1-0177. Research by the second author was also partially supported by NSF under Grant CCR-99-88216. Research by the third author was also partially supported by NSF under Grants CCR-96-19542 and CCR-97-12088. Research by the fourth author was also partially supported by NSF under Grant DMS-97-27859.

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 [17]. 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 [15]. Such physical processes are simulated through numerical computations facilitated by a mesh representing the boundary and/or domain. This mesh may be a two-dimensional triangulation of the surface, or a three-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.

We are also interested in using deforming surfaces in the modeling of molecules. Deformations happen naturally, for example in the folding process of proteins [5]. Beyond natural phenomena, we see a purpose in creating artificial deformations, for example to interpolate continuously between two time-slices in a molecular dynamics simulation, or between reconstructions of a protein from two different crystallizations.

Skin Surfaces. The approach to deforming surfaces taken in this paper is based on the technical notion of skin surfaces, as introduced in [8]. The main reason for this choice is the existence of fast combinatorial algorithms based on the theory of alpha shapes [9]. 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 obtained by shrinking the Minkowski sum of corresponding Delaunay and Voronoi polyhedra. 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 in fact can be extended naturally to a continuous function throughout space. Equally important is the knowledge about when, where, and how the surface changes its topological type. This and other geometric information is readily computable from the dual complex, the mixed complex, and the decomposition of the surface defined by the mixed complex. All this is explained shortly.

Triangulation. For computational purposes we want to approximate the skin surface by a two-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 their spacing depends on curvature. The algorithm maintains the

triangulation through local restructuring operations:

- 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 weighted points defining it. The maximum curvature at each surface point is a single real number, so our adaptation to local density produces an isotropic triangulation. We gain flexibility by permitting the triangles to deviate somewhat from the equilateral shape, and we use that flexibility for obvious geometric reasons but also for algorithmic efficiency. The deviation is measured as circumradius over length of the shortest edge, and the algorithm guarantees that this ratio never exceeds $Q^2/2$. Here Q is one of the constants on which the algorithm depends, the other being C , which controls how well the triangulation approximates the surface. Q controls how far local density can deviate from strict inverse proportionality to local maximum curvature. The two constants need to be chosen judiciously in order to guarantee the correctness of the algorithm.

Outline. The technical portion of this paper is divided into three parts and nine sections. Part I provides the geometric background. It consists of Section 2 describing skin surfaces, Section 3 showing that normal direction and maximum curvature vary slowly, and Section 4 introducing a combinatorial method for triangulating the surface. Part II explains the algorithm. It consists of Section 5 discussing adaptation to changing shape, Section 6 discussing adaptation to changing curvature, and Section 7 discussing adaptation to changing topology. Part III proves the algorithm is correct. It consists of Section 8 analyzing the adaptation to curvature, Section 9 detailing the various operations of the algorithm, and Section 10 analyzing the adaptation to changing topology. Section 11 concludes the paper with suggestions for future work.

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 [8] for the original introduction of skin surfaces, to [7] and [9] for a description of alpha shapes, and to [16] for a textbook in geometry that talks about a version of 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* $\pi_{\hat{a}}: \mathbb{R}^3 \rightarrow \mathbb{R}$ 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

$$\pi(x) = \gamma (\|x - p\|^2 - \beta),$$

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 $\alpha\hat{a} + \beta\hat{b}$ for the zero-set of $\alpha\pi_{\hat{a}} + \beta\pi_{\hat{b}}$. Using this notation, we can define what we mean by the *affine hull* and by the *convex hull* of a finite collection of spheres $\mathcal{A} = \{\hat{a}_1, \hat{a}_2, \dots, \hat{a}_n\}$, namely

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

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 $\pi_{\hat{a}} + A^2/2$. 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,

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

In other words, the skin is the boundary of the *body*, denoted body $\mathcal{A}$, which is the union of the balls bounded by spheres in $\sqrt{\text{conv } \mathcal{A}}$.

Mixed Cells. The mixed cells mentioned in the Introduction are obtained from the corresponding weighted Voronoi and Delaunay polyhedra. 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, $v_{\hat{a}} = \{x \in \mathbb{R}^3 \mid \pi_{\hat{a}}(x) \leq \pi_{\hat{b}}(x) \text{ for all } \hat{b} \in \mathcal{A}\}$. Two Voronoi polyhedra meet at most along a common piece of their boundary, and we define

$$v_{\mathcal{X}} = \bigcap_{\hat{a} \in \mathcal{X}} v_{\hat{a}}$$

for every subset $\mathcal{X} \subseteq \mathcal{A}$. It is convenient to assume general position, in which case the dimension of each non-empty $v_{\mathcal{X}}$ is $\dim v_{\mathcal{X}} = 4 - \text{card } \mathcal{X}$. In particular, $v_{\mathcal{X}}$ is a

Fig. 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.

polyhedron, polygon, edge, vertex if $\mathcal{X}$ has cardinality 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:

$$\delta_{\mathcal{X}} = \text{conv} \{a \mid \hat{a} \in \mathcal{X}\}.$$

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 $v_{\mathcal{X}}$ and $\delta_{\mathcal{X}}$ have complementary dimensions: $\dim v_{\mathcal{X}} + \dim \delta_{\mathcal{X}} = 3$. Furthermore, they lie in orthogonal affine subspaces of $\mathbb{R}^3$. We use vector operations in $\mathbb{R}^3$ to construct the *mixed cell* as a Minkowski sum,

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

The dimension of $\mu_{\mathcal{X}}$ is always 3 = $\dim v_{\mathcal{X}} + \dim \delta_{\mathcal{X}}$. 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-face tiling of $\mathbb{R}^3$, which we call the *mixed complex*. Figure 11 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 $\mathcal{X}$ [8]. Specifically, it is the same as the envelope of the affine hull after shrinking all spheres, that is,

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

Let $k = \text{card } \mathcal{X} - 1$. Then for $k = 0$ or 3 the envelope of $\sqrt{\text{aff } \mathcal{X}}$ 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* as the point $z_{\mathcal{X}}$ that is common to the affine subspaces defined by $v_{\mathcal{X}}$ and 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 $v_{\mathcal{X}} \cap \delta_{\mathcal{X}}$ is non-empty.

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 sphere and the hyperboloids are

$$x_1^2 + x_2^2 + x_3^2 = R^2, \tag{1}$$

$$x_1^2 + x_2^2 - x_3^2 = \pm R^2. \tag{2}$$

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

The plus sign gives the one-sheeted hyperboloid and the minus sign gives the two-sheeted hyperboloid. The double-cone arises as the limiting case for $R = 0$. The three surfaces are illustrated in Fig. 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 and note that the technical results in this paper are restricted to this model. It is generated by changing the original weights A^2 of the weighted points $\hat{a}$ to $A^2 + t$ at time t . It is easy to see that this weight change preserves the Voronoi polyhedra and therefore also the Delaunay simplices and mixed cells. Even though the mixed complex remains unaffected, we observe all generic types of topological changes or metamorphoses that arise in general deformations. This is illustrated in Fig. 3. As indicated in Table 1, there are four types depending on $k = \text{card } \mathcal{X} - 1$, where $\mu_{\mathcal{X}}$ is the mixed cell containing the metamorphosis. 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, we can define a collection of spheres $\mathcal{B} = \mathcal{A}^\perp$ with skin $\mathcal{A} = \text{skin } \mathcal{B}$ and body $\mathcal{A} \cup \text{body } \mathcal{B} = \mathbb{R}^3$. Essentially, $\mathcal{B}$ contains a weighted point at every Voronoi vertex $b = v_{\mathcal{X}}$, and the weight is chosen so that $B^2 = \|a - b\|^2 - A^2$ for every $\hat{a} \in \mathcal{X}$ [8]. 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 and $3 - k$. In other words, there are only two basic types of metamorphoses. The first type is geometrically realized

Fig. 3. Three snapshots of a deforming skin triangulation defined by continuously growing spheres. From left to center, note in the front two $k = 1$ metamorphoses, each adding a handle. From center to right, note at the left a $k = 2$ metamorphosis, closing a tunnel.

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

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

by a sphere appearing or disappearing. The limit configuration is a point, and in the growth model this is the center of the appearing or disappearing 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 the metamorphoses happen. We answer this in the context of the growth model by introducing certain subcomplexes of the Delaunay complex. In the literature these subcomplexes are referred to as dual or alpha complexes [9], but we use different notation and simply denote them by $K(t)$. Here $t \in \mathbb{R}$ is time as above. The growth model replaces each weight A^2 by $A^2 + t$ at time t . Restrict each Voronoi polyhedron to within the generating sphere at time t , giving

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

The complex $K(t)$ consists of all Delaunay simplices $\delta_{\mathcal{X}}$ for which the restricted Voronoi polyhedra have non-empty intersection, that is, $\bigcap_{\hat{a} \in \mathcal{X}} v_{\hat{a}}(t) \neq \emptyset$. As t increases, $K = K(t)$ grows into a progressively larger subcomplex until eventually it is the entire Delaunay complex. We sort the simplices in the order they enter the complex K . Even with assumption of general position there are ties, which we leave unresolved, by allowing more than one simplex at a given position in the ordering. The result is a sequence of collections of Delaunay simplices that captures the evolution of the complex. Every prefix of the sequence is itself a complex. Because of this property, we also have a fast algorithm for deciding how and when the homotopy type of K changes [6].

The underlying space of $K(t)$ and the body bounded by the skin at time t are homotopy equivalent [8]. 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. Assuming general position, there is a metamorphosis for every position in the ordering occupied by a single Delaunay simplex $\delta_{\mathcal{X}}$, and the type of metamorphosis is the dimension of $\delta_{\mathcal{X}}$. Whenever there are two or more simplices tied at any one position, their effects on the homotopy type of K cancel and the body does not change its topological type.

Sandwiching Spheres. We close this section by stating a rather special and important property of skin surfaces, from [8], which is 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 $S_x \in \sqrt{\text{conv } \mathcal{A}}$ and $T_x \in \sqrt{\text{conv } \mathcal{B}}$ that pass through x . These two spheres S_x and T_x are externally tangent, and have equal radius. The skin surface stays outside both spheres, and is thus tangent to them at x .

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 [8]. In a nutshell, the reason is that all spheres $S \in \text{conv } \mathcal{A}$ are orthogonal or further than orthogonal to all spheres in $T \in \text{conv } \mathcal{B}$. If we shrink S and T each to $\sqrt{2}/2$ of the original size, then the two shrunken spheres are necessarily disjoint, unless S and T are orthogonal and of the same size, in which case the two shrunken spheres touch at a point. Since the skin is the common envelope of $\sqrt{\text{conv } \mathcal{A}}$ and $\sqrt{\text{conv } \mathcal{B}}$, the two spheres passing through x must be derived from equally large spheres, which stay equally large after shrinking.

3. Continuity of Curvature

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

Maximum Curvature. Given a surface F , a point x on F , and a tangent vector $\mathbf{t}_x$, the *normal curvature* of F is that of a geodesic passing through x in the direction $\mathbf{t}_x$. The *maximum curvature* is the function $\kappa: F \rightarrow \mathbb{R}$ that maps $x \in F$ to the maximum normal curvature at x . For a hyperboloid of revolution, the minimum curvature is measured within planes containing the symmetry axis (along meridians), and the maximum curvature is measured in the orthogonal direction (along latitudes). Explicit expressions for κ are easy to compute [12, Chapter 14]. For the sphere and the hyperboloids in standard form (1) and (2), the maximum curvatures are

$$\kappa = 1/R, \tag{3}$$

$$\kappa = 1/\sqrt{\pm R^2 + 2x_3^2}, \tag{4}$$

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 (4) 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 lie on spherical shells around the origin.

Iso-curvature Lemma. *Every point $x \in \mathbb{R}^3$ belongs to exactly one hyperboloid in standard form, and the maximum curvature of that hyperboloid at x is $\kappa(x) = 1/\|x\|$.*

For either type of hyperboloid, $1/R$ is the maximum of the curvature over the whole surface. For the one-sheeted hyperboloid, R is also the radius of the smallest circle around the neck of the hourglass. For the two-sheeted hyperboloid, R is also half the smallest distance between the two sheets.

Fig. 4. The circles S_x and T_x sandwich the hyperbola. Depending on whether we revolve the hyperbola around the vertical or the horizontal axis, we get a one-sheeted or a two-sheeted hyperboloid.

Curvature Continuity. 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 $S = \sqrt{\text{conv } \mathcal{A}}$ and $T = \sqrt{\text{conv } \mathcal{A}^\perp}$. The skin of $\mathcal{A}$ is $F = \text{env } S = \text{env } T$. 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 stating that S_x and T_x have the same size. It is convenient to define $\varrho(x) = 1/\kappa(x)$, and for reasons that will become clear later we refer to $\varrho: F \rightarrow \mathbb{R}$ as the *length scale*.

Curvature Sandwich Lemma. *For every point $x \in F$, the local length scale, $\varrho(x)$, is the common radius of S_x and T_x .*

Proof. If x belongs to a sphere patch, then that patch either lies on S_x or on T_x and $\varrho(x)$ is obviously the radius. Now suppose x belongs to a hyperboloid patch. The hyperboloid is obtained by revolving a hyperbola around one of its two symmetry axes. As indicated in Fig. 4, the hyperbola is the common envelope of two families of circles, one centered along each of the two symmetry axes. By the Sandwich Property, S_x and T_x have equal radii. Because x is halfway between the centers of S_x and T_x , that radius is equal to the distance of x from the origin. By the Iso-curvature Lemma, this distance is $\|x\| = \varrho(x)$. $\square$

The sandwiching spheres, and their common radius, vary continuously with the point $x \in F$. This is easy to see for points in the interior of a sphere or hyperboloid patch, and the tangent continuity of F implies the same for points on the boundary common to two or more patches. The Curvature Sandwich Lemma thus implies that the maximum curvature varies continuously over the skin surface (except at centers, where it blows up). In fact, at every point x the local length scale $\varrho(x)$ equals the distance from x to the center z_χ of the mixed cell μ_χ that contains x .

Curvature Variation. We strengthen the result that $\kappa(x)$ is continuous by showing that it varies rather slowly. In fact, we extend its reciprocal $\varrho(x)$ to a function defined on all of $\mathbb{R}^3$ and show that $\varrho(x)$ has Lipschitz constant one. As we have seen, within any mixed cell $\mu_{\mathcal{X}}$, ϱ is simply the distance to the center $z = z_{\mathcal{X}}$. By the definition of the mixed complex, this is a continuous function on $\mathbb{R}^3$. Within $\mu_{\mathcal{X}}$, the triangle inequality gives the Lipschitz bound,

$$|\varrho(x) - \varrho(y)| = |\|x - z\| - \|y - z\|| \leq \|x - y\|.$$

By applying this to the pieces of the line segment from x to y contained in different mixed cells, we obtain the result.

Curvature Variation Lemma. *For all points x, y in space we have $|\varrho(x) - \varrho(y)| \leq \|x - y\|$.*

We note that the extension of ϱ to a function $\mathbb{R}^3 \rightarrow \mathbb{R}$ describes the length scale of all surfaces in the family defined by the 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 one-sided 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 them and to their length scales. The outward 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 speed of the angle as we rotate the point about the origin. The formulas imply that the normals of points x and y in two adjacent mixed cells are the same if x and y are mirror images of each other across the separating plane. This property is illustrated in Fig. 5.

Fig. 5. Two mixed cells with dotted circles around their centers and parallel normals of points mirrored across the separating edge. The illustration shows the case where the cells have $k = 0$ and 1.

Normal Variation Lemma. *Let x and y be points on F with distance $\|x - y\| < \varrho(x)$. The angle $\angle \mathbf{n}_x \mathbf{n}_y$ between the surface normals at x and y is at most $\arcsin \|x - y\| / \varrho(x)$.*

Proof. Consider first the case where x and y belong to the same mixed cell, and translate the coordinates so that the center is at the origin. Given x and the distance $\|x - y\|$ between the two points, the angle subtended at the origin is a maximum if $\|x\|^2 = \|x - y\|^2 + \|y\|^2$. We thus have

$$\angle \mathbf{n}_x \mathbf{n}_y \leq \arcsin \frac{\|x - y\|}{\varrho(x)},$$

as claimed.

Consider second the case where x and y lie in different mixed cells. The directed line segment from x to y passes through $i \geq 1$ planes $h_1, h_2, \dots, h_i$ separating adjacent mixed cells. Let $p_j = xy \cap h_j$ be the intersection points ordered from x to y . We construct a polygonal path that starts at x and whose length is $\|x - y\|$. It is obtained from xy by reflecting the portion after p_j across the plane h_j , for $j = i, i - 1, \dots, 1$ in this order. The endpoint y' of the path is contained inside the sphere with radius $\|x - y\|$ around x , which implies that the angle between x and y' subtended at the origin is $\varphi < \arcsin(\|x - y\| / \varrho(x))$. Since $\mathbf{n}_y$ is normal to the sphere or hyperboloid defined for the mixed cell of x that passes through y' , φ is also the angle between $\mathbf{n}_x$ and $\mathbf{n}_y$. The claim follows. $\square$

The proof of the Normal Variation Lemma does not require that x and y belong to the same skin surface. The claimed inequality holds more generally for any points $x, y \in \mathbb{R}^3$ with normals defined by the one-parameter family of skin surfaces mentioned above. Suppose the distance between x and y is $\|x - y\| < \varepsilon \varrho(x)$. Then the Normal Variation Lemma implies

$$\angle \mathbf{n}_x \mathbf{n}_y < \arcsin \varepsilon,$$

which is the form used most often in this paper.

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 $\|x - a\| < \varepsilon \varrho(x)$. The goal of this section is to prove that the restricted Delaunay triangulation defined by an ε -sampling is homeomorphic to the skin surface, provided the following condition holds:

$$(I) \quad 0 \leq \varepsilon \leq \varepsilon_0,$$

where $\varepsilon_0 = 0.279\dots$ is a root of

$$f(\varepsilon) = 2 \cos \left(\arcsin \frac{2\varepsilon}{1 - \varepsilon} + \arcsin \varepsilon \right) - \frac{2\varepsilon}{1 - \varepsilon}.$$

Note that $f(\varepsilon)$ is defined for $-1 \leq \varepsilon \leq \frac{1}{3}$, and that it is non-negative for $0 \leq \varepsilon \leq \varepsilon_0$.

Restricted Delaunay Triangulation. Let V be a finite set of points on the skin surface. We refer to these points as vertices and we denote the Voronoi polyhedron of a vertex $a \in V$ in $\mathbb{R}^3$ by v_a . The corresponding *restricted Voronoi polygon* is the intersection with the skin surface, $F \cap v_a$, which is non-empty because $a \in F$ and $a \in v_a$. The *restricted Delaunay triangulation* is the nerve of the collection of restricted polygons:

$$D_V = \left\{ \text{conv } U \mid U \subseteq V, F \cap \bigcap_{a \in U} v_a \neq \emptyset \right\}.$$

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, and 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 [2], this means the underlying space of D is homeomorphic to F . As shown in [10], it suffices to prove that every non-empty common intersection of restricted Voronoi polygons is a closed topological 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*.

We formulate this property in terms of the (unrestricted) Voronoi polyhedra. By assumption of general position, the intersection of $k + 1 = 2, 3, 4$ Voronoi polyhedra is a polygon, edge, vertex. Depending on the case, the intersection with the skin surface is to be

- case $k = 0$: a closed disk,
- case $k = 1$: empty or a closed interval,
- case $k = 2$: empty or a single point,
- case $k = 3$: empty.

The case $k = 0$ corresponds to a single Voronoi polyhedron, which has non-empty intersection with F because its generating point lies on F . We establish four technical lemmas in preparation of proving that D has the closed ball property.

Distance Claims. If two surface points lie in the same Voronoi polyhedron, then they cannot be far from each other, and if they lie on a line that is almost normal to the surface, then they cannot be close to each other. We quantify the first claim under the assumption that V is an ε -sampling.

Short Distance Claim. *If points x and y on F belong to a common Voronoi polyhedron defined by a vertex in an ε -sampling $V \subseteq F$, then $\|x - y\| < (2\varepsilon/(1 - \varepsilon))\varrho(x)$.*

Proof. Let a be the generating point of the common Voronoi polyhedron. By the ε -sampling assumption we have $\|x - a\| < \varepsilon\varrho(x)$ and $\|y - a\| < \varepsilon\varrho(y)$. Using the triangle inequality we get $\|x - y\| < \varepsilon(\varrho(x) + \varrho(y))$. The Curvature Variation Lemma now implies

$$\begin{aligned} \varrho(x) &\geq \varrho(y) - \|x - y\| \\ &> (1 - \varepsilon)\varrho(y) - \varepsilon\varrho(x), \end{aligned}$$

and hence $(1 + \varepsilon)\varrho(x) > (1 - \varepsilon)\varrho(y)$. The distance between x and y is therefore

$$\begin{aligned}\|x - y\| &< \varepsilon \left(1 + \frac{1 + \varepsilon}{1 - \varepsilon}\right) \varrho(x) \\ &= \frac{2\varepsilon}{1 - \varepsilon} \varrho(x),\end{aligned}$$

as claimed. $\square$

We get a better bound on the distance if one of the points generates the Voronoi polyhedron. Assuming $x = a$ we get $\|a - y\| < \varepsilon\varrho(y) \leq \varepsilon\varrho(a) + \varepsilon\|a - y\|$, which implies

$$\|a - y\| < \frac{\varepsilon}{1 - \varepsilon} \varrho(a).$$

We need this version of the Short Distance Claim in the proof of the Voronoi Polyhedron Lemma below.

Next we quantify the second claim, which is independent of V .

Long Distance Claim. *Suppose a line meets F in two points x and y and forms an angle smaller than ξ with the surface normal at x . Then $\|x - y\| > 2\varrho(x) \cos \xi$.*

Proof. By the Sandwich Property, there are two spheres of radius $\varrho(x)$ that both pass through x and locally sandwich the surface. The line meets the two spheres at x and at points at distance larger than $2 \cos \xi \varrho(x)$ on both sides. The skin surface contains no points inside either sandwiching sphere, which implies the claimed lower bound for $\|x - y\|$. $\square$

We play the Short and Long Distance Claims against each other and thus reach contradictions proving various claims.

Normal Lemmas. If the vertices of a short edge or a triangle with small circumcircle lie on the skin surface, then the edge or triangle lies almost flat. We quantify both claims. For an edge ab let $\mathbf{t}_{ab} = (b - a)/\|b - a\|$ be the unit tangent vector. The first result is an immediate corollary to the Long Distance Claim:

Edge Normal Lemma. *The angle between an edge ab and the surface normal at its vertex a is $\angle \mathbf{t}_{ab} \mathbf{n}_a > \pi/2 - \arcsin(\|a - b\|/2\varrho(a))$.*

A common use of the Edge Normal Lemma is when ab belongs to the restricted Delaunay triangulation of an ε -sampling. Then $\varrho(a) > (1 - \varepsilon)\varrho(x)$, where x is a point in the intersection of the dual Voronoi polygon with the skin surface. Hence $\|a - b\| < 2\varepsilon\varrho(x) < (2\varepsilon/(1 - \varepsilon))\varrho(a)$. The angle between ab and the surface normal at a is then

$$\angle \mathbf{t}_{ab} \mathbf{n}_a > \frac{\pi}{2} - \arcsin \frac{\varepsilon}{1 - \varepsilon}.$$

Fig. 6. The dashed sandwiching spheres meet the solid sphere around a in two parallel dotted circles. Vertices b and c are placed to maximize the angle between the triangle normal and the surface normal at a .

Next we consider the triangle normal lemma. We assume the angle at a inside the triangle abc is no smaller than the angles at b and c . Let R_{abc} be the radius of the circumcircle and let $\mathbf{n}_{abc}$ be the outward unit normal vector of abc .

Triangle Normal Lemma. *If a is a vertex of the triangle abc with greatest angle, then the angle between the normal of abc and the surface normal at a is $\angle \mathbf{n}_{abc} \mathbf{n}_a < \arcsin(2R_{abc}/\varrho(a))$.*

Proof. Consider the two spheres of radius $\varrho(a)$ that locally sandwich the surface at a , as shown in Fig. 6. The face angle at a is at least $\pi/3$ and the length of the edges ab and ac is at most $2R_{abc}$ each. To compute a bound on the angle between $\mathbf{n}_a$ and $\mathbf{n}_{abc}$ we assume $\|a - c\| \leq \|a - b\|$ and consider the sphere with radius $\|a - b\|$ around a . It intersects the sandwiching spheres in two parallel circles. Let $2X$ be the distance between these two circles and note that $X/\|a - b\| = \|a - b\|/2\varrho(a)$ by dropping a perpendicular from z to the midpoint of ab and using similar triangles. Hence $2X = \|a - b\|^2/\varrho(a)$. Since the angle at a is greater than or equal to the ones at b and c , bc is the longest edge of abc . The angle between the edge bc and the planes of the intersection circles is therefore less than

$$\arcsin \frac{\|a - b\|^2/\varrho(a)}{\|b - c\|} \leq \arcsin \frac{2R_{abc}}{\varrho(a)}.$$

This is an upper bound for the angle between the two normal vectors at a . $\square$

Suppose that abc belongs to the restricted Delaunay triangulation of an ε -sampling. Then $\varrho(a) > (1 - \varepsilon)\varrho(x)$, where x is a point of the intersection between the dual Voronoi edge and the skin surface. Hence $R_{abc} < \varepsilon\varrho(x) < (\varepsilon/(1 - \varepsilon))\varrho(a)$. The angle between the two normals at a is then

$$\angle \mathbf{n}_{abc} \mathbf{n}_a < \arcsin \frac{2\varepsilon}{1 - \varepsilon}.$$

Closed Ball Property. We are now ready to prove the closed ball property for the restricted Delaunay triangulation, assuming V is an ε -sampling of F satisfying Condition (I). We assume general position and consider the three cases in turn: first Voronoi edges, then Voronoi polygons, and finally Voronoi polyhedra.

Fig. 7. A Voronoi polygon intersecting the skin in a circle to the left and two intervals to the right.

Voronoi Edge Lemma. *A Voronoi edge of V intersects the skin surface in at most one point.*

Proof. Assume there is a Voronoi edge that intersects F in at least two points, x and y . Let abc be the dual triangle in the restricted Delaunay triangulation. The Triangle Normal Lemma gives an upper bound for the angle between the normal of abc and the surface normal at a . The Normal Variation Lemma gives an upper bound for the angle between the surface normals at a and x . Together they imply an upper bound for the angle ξ between $\mathbf{n}_{abc}$ and $\mathbf{n}_x$:

$$\begin{aligned}\xi &\leq \angle \mathbf{n}_{abc} \mathbf{n}_a + \angle \mathbf{n}_a \mathbf{n}_x \\ &< \arcsin \frac{2\varepsilon}{1-\varepsilon} + \arcsin \varepsilon.\end{aligned}$$

The angle ξ is also the angle between the Voronoi edge and $\mathbf{n}_x$. The Long Distance Claim implies $\|x - y\| > 2\rho(x) \cos \xi$, which by Condition (I) contradicts the upper bound $\|x - y\| < (2\varepsilon/(1-\varepsilon))\rho(x)$ implied by the Short Distance Claim. $\square$

Voronoi Polygon Lemma. *The intersection of a Voronoi polygon of V with the skin surface is either empty or a closed topological interval.*

Proof. Assume there is a Voronoi polygon whose intersection with the skin surface contains a topological circle or two topological intervals, as shown in Fig. 7. Let ab be the dual edge in the restricted Delaunay triangulation, and let x be an arbitrary point of the intersection. If x lies on a circle, then let L be the line in the plane of the polygon that intersects the circle in a right angle at x . We have $\angle L \mathbf{n}_x \leq \angle L' \mathbf{n}_x$ for any line L' in the same plane and passing through x . Choose L' to minimize the angle with $\mathbf{n}_a$. The Edge Normal Lemma implies an upper bound for the angle between L' and the surface normal at a . The Normal Variation Lemma implies an upper bound on the angle between the surface normals at a and x . Together these inequalities imply

$$\angle L \mathbf{n}_x < \arcsin \frac{\varepsilon}{1-\varepsilon} + \arcsin \varepsilon.$$

This angle is less than the upper bound for ξ in the proof of the Voronoi Edge Lemma, which implies a contradiction between the two distance claims.

In the case of two intervals let L be a line connecting x to the closest point y on the other interval. If y lies in the interior, then L intersects the interval in a right angle at y . In this case we get a contradiction with the same argument as above only with x and y interchanged. Otherwise, y is an endpoint of the interval and lies on a Voronoi edge. The angle between L and $\mathbf{n}_y$ is less than that between the Voronoi edge and $\mathbf{n}_y$. We thus get a contradiction with the same argument as used in the proof of the Voronoi Edge Lemma. $\square$

Voronoi Polyhedron Lemma. *The intersection of a Voronoi polyhedra of V with the skin surface is a closed topological disk.*

Proof. Assume there is a Voronoi polyhedron whose intersection with the skin surface contains a closed 2-manifold (without boundary), a 2-manifold with boundary other than a disk, or two disks. In the first case we let L be a line that intersects the 2-manifold in two points, x and y , and forms a right angle at x . We get a contradiction between the two distance claims as before.

For the rest of the proof, let a be the generating vertex of the Voronoi polyhedron and assume the intersection between this polyhedron and the skin surface is a 2-manifold with boundary, F' . This 2-manifold with boundary can be different from a disk either because it is non-orientable, it contains a handle, or it has at least two boundary circles. The non-orientability of F' contradicts the orientability of F . If F' has a handle but only one boundary circle, then homology theory gives us a pair of simple closed curves in F' that intersect each other transversely exactly once. Along either one of these curves, there is a point such that the line normal to F' that passes through that point meets the other curve, and hence F' again. This gives a contradiction to the two distance claims. A more elaborate argument is needed for the case where there are two or more boundary circles. Then either F' is connected, and in the simplest case is an annulus, or it is disconnected, and in the simplest case consists of two disks.

By the remark after the Short Distance Claim, the distance between a and a point $y \in F'$ is $\|a - y\| < (\varepsilon/(1 - \varepsilon))\varrho(a)$. Let L be the normal line at a and note that it contains the line segment of length $2\varrho(a)$ that connects the centers of the two spheres sandwiching the surface at a . This line segment is contained in the Voronoi polyhedron, which implies that the polyhedron is fairly tall and slim. Consider a plane that contains L and intersects at least two boundary circles of F' . Such a plane exists for else we can find a plane through L that intersects no boundary circle at all. However, then L meets F' in at least two points, and we again get a contradiction to the two distance lemmas. The plane that meets two boundary circles intersects the Voronoi polyhedron in a convex polygon and F' in at least two connected curves. One of the curves contains a . We may assume that the second curve lies on one side of L . Let L' be the line passing through its two endpoints, which both lie on the boundary of the convex polygon. The line L' intersects the sphere with radius $(\varepsilon/(1 - \varepsilon))\varrho(a)$ around a and it does not intersect the line segment connecting the centers of the two sandwiching spheres. The angle between L' and the surface normal at a is therefore $\angle \mathbf{n}_a L' \leq \arcsin(\varepsilon/(1 - \varepsilon))$. By the intermediate value theorem there is a point y on the second curve whose curve normal $\mathbf{n}'_y$ is also normal to L' . Hence $\angle \mathbf{n}_a \mathbf{n}'_y \geq \pi/2 - \arcsin(\varepsilon/(1 - \varepsilon))$. Since the surface normal at y is also normal to the tangent line parallel to L' , its angle with $\mathbf{n}_a$ is at least this large. From the Normal

Variation Lemma we get $\angle \mathbf{n}_a \mathbf{n}_y < \arcsin(\varepsilon/(1 - \varepsilon))$. Putting both inequalities together we get $\pi/2 < 2 \arcsin(\varepsilon/(1 - \varepsilon))$. This is equivalent to $\varepsilon > \sqrt{2} - 1 = 0.414 \dots$ and contradicts Condition (I). This completes the proof of the Voronoi Polyhedron Lemma for the final case where F' has at least two boundary circles. $\square$

Summary. The three Voronoi lemmas establish that for ε satisfying Condition (I), the restricted Voronoi diagram of an ε -sampling V has the closed ball property. The result of [10] 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 ε satisfying Condition (I).*

For the purpose of changing the topology of the skin surface we rely on point distributions that locally violate the ε -sampling condition. We 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 position 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

This section describes the overall algorithm and presents the details for adapting the triangulation to the changing shape of the surface. We restrict the deformation to the growth model, where the weight A^2 of every sphere is changed to $A^2 + t$ at time t . Let $t_0 < t_1$ be moments in time and let D_0, D_1 be the corresponding restricted Delaunay triangulations. The algorithm updates D_0 locally and changes it to D_1 .

Moving Vertices. The intuition for moving vertices is taken from Morse theory, which considers structures that arise in sweeping out a smooth manifold [14]. 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 $M: \mathbb{R}^3 \rightarrow \mathbb{R}$ 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 one-parameter family of diffeomorphisms from the integral lines of the gradient vector field $\text{grad } M(x)$. These diffeomorphisms $\varphi_i: F(t_0) \rightarrow F(t_i)$, with $t_i \in [t_0, t_1]$, can be composed to connect diffeomorphically any two cross sections in the time interval,

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

Fig. 8. Dotted integral lines of a solid growing circle and a solid growing hyperbola.

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 the integral lines are pieces of straight lines and hyperbolas, as illustrated in Fig. 8. To see this, note that $(2x_1, -2x_3)$ are the normal vectors of the 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 $2x_1x_3 = \pm R^2$ obtained by rotating the first family through an angle of $\pi/4$. The three-dimensional picture is obtained by revolving the hyperbolas in Fig. 8 about the x_3 -axis. The first family of hyperbolas turns into the one-parameter family of hyperboloids described by (2). The second family turns into a two-parameter family of hyperbolas each orthogonal to each of these hyperboloids.

For deformations more general than the ones in the growth model, we may not be able to determine the integral lines explicitly. Fortunately, moving vertices along integral lines is convenient but not necessary for the algorithm, and an approximation of that movement will in general suffice. For small time steps, the triangulation changes only a small amount and can be maintained with the methods described in this and the following two sections.

Parametrization. It is convenient to parametrize the integral lines by time so that points can be moved by evaluation. Each integral line is decomposed by the mixed complex into pieces of lines and hyperbolas. We first consider the case of a line inside a mixed cell constructed from a Delaunay vertex and its dual Voronoi polyhedron. After translating the center to the origin, the mixed cell is swept out by a sphere in standard form $x_1^2 + x_2^2 + x_3^2 = s$, for $s \geq 0$. We thus get integral lines that start at the origin and go to infinity, and we clip each such half-line to within the mixed cell. If the origin lies inside the mixed cell, then it is the source of an entire sphere of integral lines. We follow the usual convention and parametrize that sphere by longitude and latitude, $\theta \in [0, 2\pi)$ and $\varphi \in [-\pi, \pi]$. For each pair of angles we have a half-line $\gamma_{\theta, \varphi}: \mathbb{R}_+ \cup \{0\} \rightarrow \mathbb{R}^3$ defined by

$$\gamma_{\theta, \varphi}(s) = \begin{pmatrix} \cos \theta \cos \varphi \sqrt{s} \\ \sin \theta \cos \varphi \sqrt{s} \\ \sin \varphi \sqrt{s} \end{pmatrix}.$$

The case of a mixed cell constructed from a Delaunay tetrahedron and its dual Voronoi vertex is symmetric, with the integral lines ending rather than starting at the origin. If the origin lies inside the mixed cell, then it is the sink of an entire sphere of integral lines.

We next consider the case of a mixed cell constructed from a Delaunay edge and its dual Voronoi polygon. We assume the hyperboloid sweeping out the mixed cell is in standard form $x_1^2 + x_2^2 - x_3^2 = s$, for $s \in \mathbb{R}$. The integral lines are hyperbolas that fall and

rise along the x_3 -axis and turn a total of 90° before reaching the x_1x_2 -plane, which they approach as they disappear to infinity. We parametrize the family with the longitudinal angle, $\theta \in [0, 2\pi)$, and the minimum distance to the origin, $R \geq 0$. For each pair of parameters we get a hyperbola $\gamma_{\theta,R}: \mathbb{R} \rightarrow \mathbb{R}^3$ defined by

$$\gamma_{\theta,R}(s) = \begin{pmatrix} \pm \cos \theta \sqrt{u} \\ \pm \sin \theta \sqrt{u} \\ \pm \frac{R^2}{2\sqrt{u}} \end{pmatrix},$$

with $u = (s + \sqrt{s^2 + R^4})/2$. To check the correctness of the parametrization note that the points $\gamma_{\theta,R}(s)$ satisfy the equation of the hyperboloid and the equations of the orthogonal hyperbola. The case of a mixed cell constructed from a Delaunay triangle and its dual Voronoi edge is symmetric, with the integral lines moving in above and below the x_1x_2 -plane and turning a total of 90° before reaching the x_3 -axis, which they approach as they go to infinity.

In either case we obtain a parametrization in time by setting $s = R^2 + t$. Note that in all four cases of integral lines, the speed of the parametrization depends only on the distance to the center of the mixed cell, $\|\partial\gamma/\partial s\| = 1/(2\|\gamma(s)\|)$. This is consistent with the length of the gradient of $M(x) = \pm x_1^2 \pm x_2^2 \pm x_3^2$ being independent of the choice of signs, $\|\text{grad } M(x)\| = 2\|x\|$.

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 follow an alternative approach and take different time steps at different locations. This is done by prioritizing the four types of operations that occur at discrete moments in time, which are edge flips, edge contractions, vertex insertions, and metamorphoses. Edge flips are described below. Edge contractions and vertex insertions arise in curvature adaptation, and are described in the next section, while metamorphoses are the operations that allow topology adaptation.

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.

At any moment in time t , we consider the collection of possible next operations. Let $t_i > t$ be the time at which 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. The overall algorithm is a simple infinite loop:

loop $\tau_i = \text{NEXTOP}; D = \text{APPLY}(\tau_i)$ 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. For example, the edge of an edge flip may disappear from D . Instead of deleting these operations immediately, we

Fig. 9. Head-on view of the edge bc . The dotted line represents the skin and the shading indicates the side of the body. To the left bc is convex and to the right it is concave.

use a lazy strategy that checks an operation when it reaches the top of the priority queue.

```

Operation NEXTOP:
  repeat  $\tau = \text{EXTRACTMIN}$  until  $\text{ISOK}(\tau)$ ;
  return  $\tau$ .

```

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.

Edge Flipping. Let bc be an edge of the restricted Delaunay triangulation D at time t . It is shared by two triangles, abc and bcd . By the Voronoi Edge Lemma, the Voronoi edges dual to abc and bcd meet the skin surface in a point each. Let L_a and L_d be the lines that contain the two Voronoi edges and orient them from where they meet F towards the point $x = L_a \cap L_d$ where they cross. The point x may or may not be a Voronoi vertex. Call bc *convex* or *concave* depending on whether the dihedral angle between abc and bcd measured on the side of the body is less than or greater than π . As illustrated in Fig. 9, in the convex case the two lines pass from outside to inside the body, and in the concave case it is the other way round. Flipping the edge bc means replacing it by the other diagonal of the quadrangle bc defines. The operation can be performed unless ad is already an edge in the triangulation, in which case either b or c belongs to only three edges. The flip would then decrease that number to two edges and contradict the closed ball property of the restricted Delaunay triangulation. The three Voronoi lemmas thus imply that the flip of bc would not be attempted if ad is already in the triangulation.

```

void EDGEFLIP( $bc$ ):
  assert  $ad \notin D$ ;
  substitute  $cad, ad, adb$  for  $abc, bc, bcd$ .

```

The edge flip operation is a response to the event that the Voronoi edges dual to abc and bcd stop meeting the skin surface. This happens when x passes through F , and in this case x is necessarily a vertex of the Voronoi diagram. If bc is convex, then x passes from inside to outside the body and ad is concave. Symmetrically, if bc is concave, then x passes from outside to inside the body and ad is convex. The time t_i when the flip happens depends on the points a, b, c, d and the surface F , all of which move continuously with

time. In other words, t_i is a root of a continuous function in t . It is used as the priority of the edge flip τ_i stored in the event scheduling priority queue.

6. Curvature Adaptation

This section focuses on the density adaptation algorithm, implemented through edge contraction and vertex insertion. 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 triangulate locally 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 R_{abc} of the circumcircle, which length scale they should follow exactly. 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 Size Bound and the Upper Size Bound:

- [L] $R_{ab}/\varrho_{ab} > C/Q$ for every edge $ab \in D$.
- [U] $R_{abc}/\varrho_{abc} < CQ$ for every triangle $abc \in D$.

It is not necessary to check for long edges and small triangles explicitly. This is because an edge of size $R_{ab} \geq CQ\varrho_{ab}$ belongs to two triangles that both violate [U]. Symmetrically, a triangle of size $R_{abc} \leq (C/Q)\varrho_{abc}$ has three edges that violate [L]. Appropriate values of C , Q will be determined in the analysis of the algorithm but we can already anticipate $C = 0.08$, $Q = 1.65$ as a feasible assignment.

Minimum Angle. The smallest angle is a measure of triangle quality. It achieves its maximum, $\pi/3$, for the equilateral triangle. 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 $\arcsin(1/Q^2)$.*

Proof. Let abc be a triangle, with bc its shortest edge and R its circumradius. We get $\varrho_{abc}/\varrho_{bc} \leq 1$ by definition of length scale. Using [L] and [U] we get

$$\frac{R}{\|b - c\|} < \frac{CQ\varrho_{abc}}{2(C/Q)\varrho_{bc}} \leq \frac{Q^2}{2}.$$

The minimum angle is $\angle bac$, and $\|b - c\| = 2R \sin \angle bac$. Hence $\angle bac = \arcsin(\|b - c\|/2R) \geq \arcsin(1/Q^2)$, as claimed. $\square$

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 $21.54 \dots^\circ$, and the maximum angle is smaller than $180^\circ - 2 \cdot 21.54^\circ = 136.90 \dots^\circ$.

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, $R_{abc} = C Q \varrho_{abc}$ at time t_i . The time t_i depends on the points a, b, c and their length scales $\varrho(a), \varrho(b), \varrho(c)$, which all change continuously with time. 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  $\exists$  triangle  $abc$  violating [U] do
    ADD( $x, abc$ )
  endwhile.
```

The details of the algorithm for adding x are discussed below.

Consider next an edge ab that gets too short, that is, $R_{ab} = (C/Q)\varrho_{ab}$ at time t_j . The moment in time t_j depends on points a, b and their length scales $\varrho(a), \varrho(b)$, which all change continuously with time. To remedy the violation of the Lower Size Bound, we contract ab by removing the vertex b with larger length scale 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  $\exists$  edge  $ab$  violating [L] do
    if  $\varrho(a) > \varrho(b)$  then  $a \leftrightarrow b$  endif;
    REMOVE( $b$ ); VERTEXINSERTION
  endwhile.
```

We note that edge contractions and vertex insertions can also be used to modify the restricted Delaunay triangulation of an oversampling until it satisfies both Size Bounds. This is because every oversampling is also an ε -sampling and thus the General Homeomorphism Theorem applies. However, if we start with an undersampling, then the algorithm may fail because conditions needed for its correct operation can be violated.

Vertex Insertion. Let L be the line of points at equal distance from a, b, c . It intersects the plane of abc in the circumcenter z of the triangle, and it intersects F in two or more points of which we add the point $x \in L \cap F$ closest to z . We prove in Section 8 that the distance between x and z is necessarily small compared with ϱ_{abc} . However, even though x and z are rather close, they may lie in different mixed cells. To find x , we start at z , determining the mixed cell that contains it by walking from a . Then we walk along

L to at most the distance specified in the Circumcenter Lemma in both directions from z . Each step in the walk enters a new mixed cell μ . Let y be the point closest to z that belongs both to L and the sphere or hyperboloid of μ . If y lies outside μ the search continues. Otherwise, $x = y$ and we add x to the triangulation by connecting it to the edges of abc .

Adding x to D as described is likely to compromise the dual correspondence to the restricted Voronoi diagram. Inspired by the incremental algorithm for constructing Delaunay triangulations in $\mathbb{R}^2$ [13], we repair the correspondence by edge flipping. More specifically, we push each edge in the link of x on an empty stack and then process the stack until it runs empty. Let pq be the top edge on the stack shared by triangles xpq inside and pqy outside the star of x . Depending on a local geometry test, we either leave pq as an edge in the triangulation or we flip it as described in Section 5. In the latter case we push the new link edges py, yq on the stack.

```
void ADD( $x, abc$ ):
  assert INSPHERE( $x, abc$ );
  substitute  $xab, xb, xbc, xc, xca, xa, x$  for  $abc$ ;
  PUSH3( $ab, bc, ca$ );
  while not ISEMPY do
     $pq = \text{POP}$ ;
    if INSPHERE( $x, pqy$ ) then
      EDGEFLIP( $pq$ ); PUSH2( $py, yq$ )
    endif
  endwhile.
```

Let z be the dual vertex of pqy in D before x was inserted. Function INSPHERE decides whether or not x lies inside the sphere with center z that passes through p, q, y . If it does, then pqy loses the reason for its existence which justifies the flip.

Edge Contraction. We contract an edge ab by removing b from the triangulation, as illustrated in Fig. 10. The operation removes b together with all edges and triangles in its star, and it covers the thus created hole by a triangulation without interior vertices. The boundary of the hole is a topological circle, which we refer to as a polygon. For each vertex p let p^- and p^+ be its predecessor and successor in an ordering around the polygon. The algorithm converts the star of b into the new triangulation by creating a triangle p^-pp^+ at a time by flipping. When only three triangles remain in the star of b

Fig. 10. The removal of b replaces the star of b by the dotted polygon triangulation.

we replace its star by a single triangle. To get started, we push all vertices in the link of b on an empty stack.

```

void REMOVE( $b$ ):
  for all  $p$  in link of  $b$  do PUSH( $p$ ) endfor;
  while stack contains more than three vertices do
     $p = \text{POP}$ ;
    if IND( $p^- p p^+$ ) then
      EDGEFLIP( $bp$ ); drop  $p$  from polygon;
      if  $p^-$  not on stack then PUSH( $p^-$ ) endif;
      if  $p^+$  not on stack then PUSH( $p^+$ ) endif
    endif
  endwhile;
   $p, q, r = \text{POP}^3$ ;
  substitute  $pqr$  for  $bpq, bq, bqr, br, brp, bp, b$ .

```

Function IND returns true iff all other vertices of the current polygon lie outside the circumsphere of p^-, p, p^+ whose center is the dual restricted Voronoi vertex.

7. Topology Adaptation

The way the skin surface is connected can change during deformation. This section studies when, where, and how these changes happen in the growth model. 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. We recall that the growth model of deformation is 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

$$x_1^2 + x_2^2 \pm x_3^2 = \pm R^2 + \frac{t}{2}.$$

Compare this with (1) and (2) in Section 2. A metamorphosis happens when the right-hand side vanishes at time $t = \mp 2R^2$, 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. We will see that in this case the metamorphosis does not happen, but we still have to modify the sampling strategy because the curvature grows beyond any bound.

Using local considerations, we can reduce the list of metamorphoses to the four given in Table 1, Section 2. 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 $M: \mathbb{R}^3 \rightarrow \mathbb{R}$ 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. We analyze the situation in some detail. 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,

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

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 $\mathcal{X}$. A hot ball 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 $\mathbb{R}_H^3$.

Hot Spot Lemma. $F_H = F \cap \mathbb{R}_H^3$.

In words, a point $x \in F$ belongs to the hot portion of the skin surface iff it belongs to the hot portion of space. In the growth model the hot portion of space is constant, while the hot portion of the skin changes as the surface moves through that portion of space. The Hot Spot Lemma follows directly from the Iso-curvature Lemma and does not need a separate proof.

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 Fig. 11. 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

Fig. 11. Dotted Voronoi diagram, dashed Delaunay triangulation, solid mixed complex, solid data points, hollow other centers, and shaded hot portion of space. Each label shows the dimension of the Delaunay simplex involved in the construction of the mixed cell.

shortly discover that an even stronger separation property between hot balls is needed to prevent edges reaching from one 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. For now, we simplify the discussion by assuming the non-degenerate case, where H is small enough such that hot balls are pairwise disjoint. We will return to the degenerate case shortly.

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 an arbitrarily small size. The case of a hyperboloid that approaches its limiting double-cone is more problematic, because the number of vertices near the center grows beyond any bound. To circumvent the computational impossibility of sampling infinitely many points, we change the sampling strategy 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 to the changed sampling density and keeps constructing the restricted Delaunay triangulation.

Consider a two-sheeted hyperboloid and translate time such that the metamorphosis happens at time $t = 0$. The hyperboloid enters its hot ball at time $-2H^2$, turns into a double-cone at time 0, and leaves the hot ball as a one-sheeted hyperboloid at time $2H^2$. The special sampling strategy that allows us to go through this motion depends on a parameter $0 < h < 1$. Special sampling begins at time $t_0 = -2H^2h^2$ when the two-sheeted hyperboloid enters the ball of radius Hh , and it ends at time $t_1 = 2H^2h^2$ when the one-sheeted hyperboloid leaves that ball, as shown in Fig. 12. At time t_0 , the hyperboloid intersects the boundary of the hot ball 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* to pass through the two hot circles. At time 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.

Fig. 12. Head-on view of start, middle, end configurations generated by special sampling taking a two-sheeted to a one-sheeted hyperboloid. The hot sphere is solid and the sphere that triggers the metamorphosis is dotted.

Extreme Configuration. At time t_0 , we kick off special sampling by creating the *double-cup* as the start configuration of the metamorphosis representing the intruding portion of the two-sheeted hyperboloid, as shown in the left drawing of Fig. 12. Consider one sheet of the hyperboloid and let a be its intersection point with the symmetry axis. Let $b_0, b_1, \dots, b_{\ell-1}$ be the vertices of a regular ℓ -gon along the hot circle in this sheet. We mirror these points across the symmetry plane of the hyperboloid and get points $a', b'_0, b'_1, \dots$ on the other sheet.

```

ADD2( $a, a'$ );
for  $i = 0$  to  $\ell - 1$  do ADD2( $b_i, b'_i$ ) endfor;
EDGECONTRACTION.

```

Recall that Function ADD really takes two parameters, namely a point and a triangle whose circumsphere centered at the dual restricted Voronoi vertex encloses the point. Each call to the function must therefore be preceded by a search for such a triangle. Whenever we contract an edge by removing one of its endpoints we make sure that the endpoint is not one of the newly added vertices. Section 10 derives sufficient conditions for h and ℓ that guarantee the above algorithm successfully constructs the double-cup as the start configuration of the metamorphosis. By this we mean that

- (i) a, a' are the only vertices inside and b_i, b'_i , for $0 \leq i < \ell$, are the only vertices on the hot sphere,
- (ii) the link of a in D is the regular ℓ -gon of vertices b_i , and symmetrically the link of a' is the ℓ -gon of vertices b'_i .

Assuming $C = 0.08$, $Q = 1.65$ we will see that $h = 0.98$ and $\ell = 5$ are feasible values for the two constants. For ease of reference we say the vertices and edges in the links of a, a' are *hot* and the vertices, edges, and triangles in the stars of a, a' are *very hot*.

The end configuration of the metamorphosis is similar to the start configuration of the inverse metamorphosis. As shown in the right drawing of Fig. 12, it consists of two rings of triangles forming a *cylinder-with-a-waist* representing the intruding portion of the one-sheeted hyperboloid. Let $u_0, u_1, \dots, u_{m-1}$ be the vertices of a regular m -gon along the waist where the hyperboloid intersects its symmetry plane. Similarly, let $w_0, w_1, \dots, w_{m-1}$ be the vertices of another regular m -gon along one of the two hot circles, rotated by π/m relative to the m -gon along the waist. Finally, let $w'_0, w'_1, \dots, w'_{m-1}$ be the vertices of the mirror m -gon on the other hot circle.

```

for  $j = 0$  to  $m - 1$  do ADD3( $u_j, w_j, w'_j$ ) endfor;
EDGECONTRACTION.

```

As before we search for an offending triangle before we add a point, and we make sure that the removed vertex of every contracted edge pq is not one of the newly added ones. Section 10 derives sufficient conditions for h and m that guarantee the above algorithm successfully constructs the cylinder-with-a-waist as the start configuration of the inverse metamorphosis. What precisely we mean by this should be obvious. Assuming $C = 0.08$, $Q = 1.65$ we will see that $h = 0.98$ and $m = 40$ are feasible values for the two constants. For ease of reference we again say that the vertices and edges along the

two hot circles are *hot* and that the vertices, edges, and triangles between the two hot circles are *very hot*.

In the forward direction we switch from the double-cup to the cylinder-with-a-waist at time 0, and in the backward direction we do it the other way round. The latter is easier because we just need to meld the m -gon of the waist into a single vertex and then split that vertex into two. In the forward direction we first meld a and a' into a single vertex and then expand that vertex into a regular polygon (interleaving angularly between the b_i). The expansion creates two new rings of triangles between the new polygon and the polygons representing the two hot circles. This is done following the angular order of the involved vertices around the symmetry axis.

Special Sampling. The main difference between special and general sampling is that the former gives up on the Lower Size Bound for hot edges and on the Upper Size Bound for very hot triangles. The length of hot edges is bounded from above because the Upper Size Bound applies to the incident triangles outside the hot sphere. A more detailed analysis of edge and triangle sizes including a proof of the closed ball property in spite of special sampling is given in Section 10.

The goal of special sampling is to maintain the double-cup and the cylinder-with-a-waist during the first and the second halves of the time interval. It acts primarily by modifying general sampling for points on and inside the hot sphere. As a general rule, an edge is contracted by removing an endpoint that is not hot. Cases where both endpoints are hot occur only at the end of the metamorphosis (or its inverse) and will be discussed separately. There are two ways in which general sampling can intrude into the hot sphere: by adding a point inside that sphere and by flipping a hot edge. In both cases we prevent the intrusion by bisecting the endangered hot edge bc , as illustrated in Fig. 13. Specifically, we add the midpoint q of the shorter hot circle arc that connects b with c . The addition of q may create edges that violate the Lower Size Bound. Of these we contract the ones that are not hot, always making sure we remove the endpoint that is not hot. As discussed above, we choose H small enough so that hot spheres cannot get too close to each other and every non-hot edge has at least one non-hot vertex. Infinite loops cannot occur because each iteration leaves an additional hot vertex behind. The hot circle gets denser and intrusions into the hot sphere get progressively more difficult.

Special sampling maintains the special configuration, but it does not guarantee the two Size Bounds. They must therefore be enforced algorithmically at the end of the metamorphosis.

SPECIAL VERTEX INSERTION;
EDGE CONTRACTION.

Fig. 13. The hot edge bc is bisected either because the dual restricted Voronoi vertex of bcd lies inside the hot sphere or edge flipping attempts to change bc to da .

The clean-up operation is correct if we maintain the closed ball property, which is initially guaranteed by special sampling. While maintaining that property might be difficult in general, we can use the insights gained from the proofs of the two Persistence Lemmas in Section 10 and add points only on the two hot circles. This is the difference between the above function and Function VERTEXINSERTION introduced earlier. The size analysis in Section 10 implies that we can satisfy the Upper Size Bound even with this restriction on new vertex locations.

Degenerate Centers. Recall that a degenerate center is one that lies on the boundary of its mixed cell. Each facet lies half-way between the centers of the two mixed cells that share it. If it contains one center, then it also contains the other, which implies that a degenerate center is also a multiple center. In a non-degenerate mixed complex (which may still have degenerate centers) every facet is shared by two, every edge by four, and every vertex by eight mixed cells. In the Morse theoretic view of centers as critical points, each degenerate center is the location where several critical points collide and cancel each other. This is why there is no metamorphosis at those locations. When the surface moves through such a degenerate center, its curvature does blow up momentarily, but the surface then becomes smooth again, with no topology change.

The remainder of this section describes the various types of degenerate centers. The enumeration exhausts the cases that occur in non-degenerate mixed complexes. Consistent with Table 1, the label of a mixed cell is k if it is constructed from a k -dimensional Delaunay simplex and its dual $(3 - k)$ -dimensional Voronoi polyhedron. We label the facets, edges, and vertices by concatenating the labels of the mixed cells that share them. There are three facet types labeled 01, 12, 23, two edge types labeled 0112, 1223, and one vertex type labeled 01112223.

The case where the degenerate center lies in the interior of a 01 facet is illustrated in Fig. 14. Reading the top row of the figure from left to right we see the skin surface passing through the center. The bottom row shows the hot ball around the degenerate center and its intersection with the evolving surface and the body it bounds. The case of a degenerate center on a 23 facet is symmetric. The shrunken Voronoi polygon is replaced

Fig. 14. Degenerate center in the interior of a 01 facet. Evolution of skin surface at the top and of hot ball at the bottom. The shaded regions on the hot spheres show the intersection with the body.

Fig. 15. Degenerate center in the interior of a 12 facet. Evolution of skin surface at the top and of hot sphere at the bottom. The hyperboloid in the type-2 cell has a vertical symmetry axis, while the hyperboloid in the type-1 cell behind it has a horizontal one.

by a shrunk Delaunay triangle, and we read Fig. 14 from right to left. Furthermore, the body lies on the other side of the skin surface.

The case of a degenerate center in the interior of a 12 facet is different and geometrically more interesting. Both mixed cells contain hyperboloids, and their symmetry axes pass orthogonally through the common center. Each symmetry axis lies in the symmetry plane of the other hyperboloid, which forms a right angle with the plane of the facet. The evolution of the skin surface passing through such a degenerate center is illustrated in the top row of Fig. 15. As shown in the bottom row, the skin surface intersects the boundary of the hot ball around the degenerate center in a curve consisting of four half-circles. Instead of a pair of hot circles, we have a hot curve consisting of four half-circles.

The case of a degenerate center in the interior of an edge labeled 0112 is illustrated in Fig. 16. The two pairs of facets common to a type 1 cell form a right dihedral angle each, and the remaining two dihedral angles add up to π but are otherwise arbitrary. As shown to the right, the skin surface intersects the boundary of the hot ball in four circular arcs, two of which are half-circles. The angles remain the same during the entire transition

Fig. 16. Degenerate center in the interior of a 0112 edge. Snapshot of skin surface and hot ball at the moment the surface passes through the center.

in which the skin surface sweeps through the hot ball. The case of a degenerate center on an edge labeled 1223 is symmetric, and Fig. 16 applies, except that the body lies on the other side of the surface. In the case of a degenerate center at a vertex, the hot curve consists of three pairs of circular arcs. The angles of the arcs remain the same during the transition in which the skin surface sweeps through the hot ball.

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 scheduling, and Section 10 examines the topology adaptation algorithm.

8. Sampling Density

We derive conditions for the constants C and Q 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 is not too large. We also prove that the restricted Voronoi vertex dual to a triangle can be found near the circumcenter of that triangle. That proof requires that the vertices of the triangulation form an ε -sampling. Finally, we 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 triangulation produced 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.

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

$$(III) \quad \delta^2/(1 + \delta)^2 - \delta^4/4 > C^2Q^2,$$

where $\delta = \varepsilon - 2C(\varepsilon + 1)/(Q + 2C)$. We get (II) and (III) as sufficient conditions for the proofs of the No-Short-Edge Lemma and the Sampling Lemma below. Condition (II) is equivalent to $Q > 2C + \sqrt{4C^2 + 2}$. Assuming $\varepsilon = \varepsilon_0 = 0.279\dots$, we can satisfy Conditions (II) and (III) by setting $C = 0.08$ and $Q = 1.65$. In this case $\delta = 0.166\dots$. 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 create edges of size below the allowed threshold. This is what prevents infinite loops in spite of the algorithm's partially conflicting efforts to avoid short edges and large triangles simultaneously. 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 $(C/Q)\varrho_{xy}$.*

Proof. We have $R_{abc} \geq CQ\varrho_{abc}$. The sphere with center x that passes through a, b, c has radius $X \geq R_{abc}$ and it contains no vertex other than x inside. Every new edge xy has therefore length $\|x - y\| \geq X \geq CQ\varrho_{abc}$. Assume without loss of generality that $\varrho_{abc} = \varrho(a)$. We use the Curvature Variation Lemma to derive upper bounds for the length scales at x and y :

$$\begin{aligned}\varrho(x) &\leq \varrho(a) + X \leq \left(\frac{1}{CQ} + 1\right) \|x - y\|, \\ \varrho(y) &\leq \varrho(x) + \|x - y\| \leq \left(\frac{1}{CQ} + 2\right) \|x - y\|.\end{aligned}$$

Hence

$$R_{xy} = \frac{\|x - y\|}{2} \geq \frac{\max\{\varrho(x), \varrho(y)\}}{4 + 2/CQ}.$$

Condition (II) implies $C/Q < CQ/(4CQ + 2)$, and therefore $R_{xy} > (C/Q)\varrho_{xy}$, as claimed. $\square$

Close Dual Vertices. Consider the point addition triggered by the triangle abc violating the Upper Size Bound. As before, we denote the line of points at equal distance from a, b, c by L , the circumcenter of abc by z , and the point of $L \cap F$ closest to z by x . We prove an upper bound on the distance between x and z assuming an ε -sampling of F .

Circumcenter Lemma. *The distance between x and z is $\|x - z\| < (\varepsilon^2/2)\varrho_{abc}$.*

Proof. Assume $\varrho_{abc} = \varrho(a) \leq \varrho(b), \varrho(c)$. We have $\|x - a\| \leq \varepsilon\varrho(x)$ by assumption of ε -sampling and therefore $(1/(1 + \varepsilon))\varrho_{abc} \leq \varrho(x)$ by the Curvature Variation Lemma. We get an upper bound on the distance between x and z by assuming $\varrho(x)$ is as small as possible and a, b, c lie on the sandwiching sphere with radius $\varrho(x) = (1/(1 + \varepsilon))\varrho_{abc}$ passing through x . This configuration is sketched in Fig. 17. Note that $\|x - z\|/\|x - a\| = \|x - a\|/2\varrho(x)$ by equality of angles formed by orthogonal sides. Therefore,

$$\|x - z\| = \frac{\|x - a\|^2}{2\varrho(x)} \leq \frac{\varepsilon^2}{2}\varrho(x),$$

as claimed. $\square$

Fig. 17. Dashed sphere of radius $\varrho(x)$ passing through a, b, c, x and bold circle with center z passing through a, b, c .

The relevance of the Circumcenter Lemma to the curvature adaptation algorithm should be obvious. When the triangle abc violates Condition (III), we need first to find its dual vertex in the restricted Voronoi diagram and then add this vertex to V . This vertex is the point x , and the Circumcenter Lemma gives a bound on how far from z we have to search before we are guaranteed to find x . As shown in the proof of the Voronoi Edge Lemma, each additional point $y \in L \cap F$ is too far from z to possibly belong to the Voronoi edge dual to abc .

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 $a \in V$ whose distance from x is $\|a - x\| < \varepsilon \varrho(x)$. 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 each other on 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 of the way it would temporarily have to violate 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. For now we disallow metamorphoses. Let $t_0 < t_1$ be two points in time so the topological type is constant within $[t_0, t_1]$. 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 $F(t_1)$.*

Proof. Assume the opposite and let $t \in [t_0, t_1]$ be the first moment in time when the skin surface is not ε -sampled. Then there is a point $x \in F(t)$ such that no vertex lies inside the sphere with center x and radius $X = \varepsilon \varrho(x)$. By minimality of t , the sphere passes through at least one vertex, a , but we need three. To get two more, we continuously increase the sphere while keeping its center on the surface. Vertex a remains on the sphere at all times and we permit no vertices inside the sphere. Let $y \in F(t)$ be the center when we reach the other two vertices, b and c . The radius of the new sphere is $Y \geq X$ because the radius can only increase from x to y . Using the Curvature Variation Lemma, we get $\varrho_{abc} \leq \varrho(x)(1 + \varepsilon)$ and therefore $(\varepsilon/(1 + \varepsilon))\varrho_{abc} \leq \varepsilon \varrho(x) \leq Y$. Assume without loss of generality that $\varrho_{abc} = \varrho(a)$, and let z be the circumcenter of abc . The Upper Size Bound implies $\|z - a\| = R_{abc} < CQ\varrho(a)$. Using the Circumcenter Lemma, we get an upper bound on the square distance between y and a ,

$$\begin{aligned} Y^2 &= \|y - z\|^2 + \|z - a\|^2 \\ &< \frac{\varepsilon^4}{4} \varrho_{abc}^2 + C^2 Q^2 \varrho_{abc}^2. \end{aligned}$$

This implies

$$\frac{\varepsilon^2}{(1 + \varepsilon)^2} < \frac{\varepsilon^4}{4} + C^2 Q^2,$$

which contradicts Condition (III). $\square$

For example for $C = 0.08$, $Q = 1.65$ the Sampling Lemma holds for all ε in an interval with endpoints $0.15 \dots$ and $0.98 \dots$.

The algorithm explained in Section 7 maintains an ε -sampling across metamorphoses. More precisely, it violates the required sampling density inside the hot sphere of each metamorphosis. As proved in Section 10, the special sampling strategy repairs the ε -sampling property before the skin surface comes out of the hot sphere, and it maintains the closed ball property at all times. The Sampling Lemma thus generalizes to any time interval, including those that contain metamorphoses.

Denser Sampling. Even though the Sampling Lemma proves that the vertices form an ε -sampling between all operations, it allows for momentary violations of the density requirement during the contraction of an edge ab . Specifically, ε -samplings are not guaranteed right after b is removed and before appropriate vertex insertions repair the Upper Size Bound. The contraction happens only if the Lower Size Bound is violated, which implies $\|b - a\| \leq (2C/(Q))\varrho_{ab}$. Since such an edge contraction always removes the vertex with larger length scale we have $\varrho(b) \geq \varrho(a)$ and therefore $\varrho_{ab} \leq \varrho(x) + \|x - b\|$. For every point $x \in F$ whose only vertex within distance $\varepsilon\varrho(x)$ is b we thus have

$$\begin{aligned} \|x - a\| &\leq \|x - b\| + \|b - a\| \\ &\leq \varepsilon\varrho(x) + \frac{2C}{Q}(\varrho(x) + \varepsilon\varrho(x)) \\ &= \left[\varepsilon + \frac{2C}{Q}(\varepsilon + 1) \right] \varrho(x). \end{aligned}$$

In words, the vertices still form an ε -sampling, but for a somewhat larger value of ε . To say the same thing in reverse we define $\delta = \varepsilon - (2C/(Q + 2C))(\varepsilon + 1)$ and note that $\varepsilon = \delta + (2C/Q)(\delta + 1)$. If we choose the constants C and Q such that the Sampling Lemma implies the maintenance of a δ -sampling, then we have an ε -sampling at all times, even during the execution of an edge contraction. Condition (III) enforces such a choice of constants.

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 others 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. It also provides correctness proofs for the flipping algorithms used to restore the restricted Delaunay triangulation after vertex insertions and edge contractions.

Total and Partial Ordering. Operations triggered by the motion of the skin surface are ordered in time. We have five types:

1. coordinate updates,
2. edge flips,
3. edge contractions,

4. vertex insertions,
5. metamorphoses.

Vertex coordinates change continuously with time, and we avoid most of the related computational expense by updating coordinates 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, or 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 is also 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. For example a vertex insertion relies on point additions and edge flips to achieve the desired effect locally and restore the restricted Delaunay triangulation. Conceptually, such a sequence is executed at an instant, while time stands still. We cannot therefore 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 for point removal described in Section 6.

Flipping after Point Addition. A vertex insertion operation is recursive and unwinds into a sequence of point additions (usually just one), each followed by a sequence of edge flips. Let D_0 and D_1 be the restricted Delaunay triangulations immediately before and after adding the vertex x . We focus on the sequence of flips following the addition of x and argue that this sequence successfully constructs D_1 .

As explained in Section 6, the algorithm maintains a subset of the edges in the link of x on a stack. Each edge pq in the link belongs to two triangles, xpq inside and pqy outside the star of x . Thus pqy exists in D_0 and it remains in D_1 iff x lies outside the circumscribed sphere of pqy whose center is the dual restricted Voronoi vertex. The INSPHERE test used to decide whether or not to flip pq captures exactly that information. If the decision is negative we know that pqy and pq remain in D_1 . Otherwise, the flip of pq increases the star of x by one triangle, and it decreases the portion of the triangulation outside the star by one triangle. The monotonicity of the transfer implies that the loop of edge flips halts. There are no obstacles to flipping pq other than if xy is already an edge of the triangulation, which is the case iff either p or q belongs to only three edges. Say p belongs to px , pq , py and to no other edges. By the closed ball property of the restricted Voronoi diagram, p belongs to at least three edges of D_1 , and because p can only lose edges during flipping, the INSPHERE test for x , pqy must be negative and the flip of pq will not be attempted. Another conceivable reason for failure is that the link of x does not reach a triangle that should be removed. However, this is impossible because the closed ball property implies that $D_0 - D_1$ is an open disk, and the link of x , which is a topological circle, is adjacent to triangles in $D_0 - D_1$ until it has swept out that entire disk.

Flipping after Point Removal. An edge contraction unwinds into a sequence of point removals and point additions, each followed by a sequence of edge flips. Let D_0 and D_1 be the restricted Delaunay triangulations immediately before and after removing a point b . We argue that the sequence of edge flips following the removal successfully constructs D_1 .

As explained in Section 6, only the polygon bounded by the link of b requires retriangulation. The algorithm flips one diagonal and recurses for the remaining star of b until only three triangles remain. It thus halts after a number of edge flips that is less than the number of triangles in the initial star. To see that the algorithm is correct, we observe that each flip generates a triangle that is guaranteed to belong to D_1 . The membership in D_1 is guaranteed by Function IND, which checks all remaining vertices of the polygon and not just one as for the flips following a point addition. Edge flips that cannot be executed because one of the endpoints has degree 3 will again not be attempted because they contradict the closed ball property of the restricted Voronoi diagram.

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 two-sheeted hyperboloid that enters the ball with radius Hh around its center, and we end with a one-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 Fig. 12. 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 i , the lengths of the edges of abc are

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

Any isosceles triangle with sides of length E and height L has circumradius $E^2/2L$. The height of abc is $L_{a,bc} = \sqrt{4R_{ab}^2 - R_{bc}^2}$. The circumradius is therefore $4R_{ab}^2/2L_{a,bc}$, which is

$$R_{abc} = \frac{R_1(R_1 - Hh)}{\sqrt{2R_1(R_1 - Hh) - R_0^2 \sin^2(\pi/\ell)}}.$$

Next refer to the cylinder-with-a-waist shown in Fig. 12. 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 i , the lengths of the edges uv and wx are

$$2R_{uv} = 2Hh \sin \frac{\pi}{m},$$

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

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 $R_{ab}, R_{bc}, R_{uw}/h, R_{wx}, R_{vw} > (C/Q)H$ and $R_{abc}, R_{uvw}, R_{vwx} < CQHh$. The inequalities for R_{vw}, R_{uw}, R_{uvw} are redundant because $R_{ab} < R_{vw}, R_{wx} < R_{uw}/h, R_{uvw} < R_{vwx}$ for all $h < 1$. In addition to requiring that the triangles abc, uvw, vwx satisfy the Upper Size Bound, it is convenient to require

also that their radii are less than the locally allowed minimum edge length. This extra requirement implies that after adding points on and inside the hot sphere, all old points inside or on the hot sphere are too close to at least one new point and thus get deleted. It follows that all remaining old vertices lie outside the hot sphere. We thus have the following two conditions:

- (IV) $R_{ab}/H, R_{bc}/H, R_{wx}/H > C/Q$,
 (V) $R_{abc}/H, R_{vwx}/H < \min\{Q, 2/Q\}Ch$.

Conditions (I)–(V) are satisfied for $\varepsilon = 0.279$, $C = 0.08$, $Q = 1.65$, $h = 0.98$, $\ell = 5$, $m = 40$. 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 $(2/Q)Ch$ 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, which follows from the Persistence Lemmas proved below.

Persistence of Triangulation. We show that the special configurations exist as sub-complexes of the restricted Delaunay triangulation during the entire time interval of a metamorphosis. Consider the simplices in the Delaunay complex spanned by hot vertices and their dual Voronoi polyhedra. Fig. 19 sketches both for the double-cup before and the cylinder-with-a-waist after the middle of the time interval. During the first half of the time interval, the hot vertices span two pyramids, one being the reflection of the other across the symmetry plane of the hyperboloid. The points are in degenerate position, which implies that the Delaunay complex Δ_1 of the hot points contains polyhedra that are more com-

Fig. 19. Solid restricted Delaunay triangulations and dotted Voronoi polyhedra of hot ball configuration before and after the double-cone.

plicated than tetrahedra. Specifically, Δ_1 consists of the two pyramids joined by an edge connecting their apices and a ring of five-sided polyhedra and quadrangles around that edge. As usual, D denotes the restricted Delaunay triangulation of the entire vertex set V .

Persistence Lemma A. *At any time in $[-2H^2h^2, 0)$, the intersection of Δ_1 and D consists of the two rings of triangles forming the double-cup and of their edges and vertices.*

Proof. We first show that the edges, polygons, and polyhedra in Δ_1 that do not belong to the double-cup also do not belong to D . The edges connecting the two cups have dual Voronoi polygons which lie in the symmetry plane separating the two sheets and therefore cannot intersect the hyperboloid. To see that they do not intersect any other part of F , we consider the sandwiching spheres defined for points of F inside the hot sphere. The Voronoi polygons are contained in the union of balls bounded by these spheres, else they would imply an empty sphere that intersects the hyperboloid in a patch outside the hot sphere that is large enough to contradict the ε -sampling property. Detailed computations of a lower bound for the size of such an implied patch are omitted. Since D is a complex, it also does not contain the Delaunay polygons and polyhedra incident to the excluded edges. The base polygons of the two pyramids in Δ_1 have their dual Voronoi edges on the symmetry line of the hyperboloid. For the same reason as above, these edges are contained in the union of balls bounded by the sandwiching spheres of points of F inside the hot sphere.

We second show that the triangles abc of the double-cup belong to D . At time $t_0 = -2H^2h^2$ this is true because these triangles have circumspheres that are small enough that every point of F inside these spheres would belong to edges that violate the Lower Size Bound. At times $t_0 < t < 0$ this is true because any violation is prevented by the algorithm before it occurs. $\square$

During the second half of the time interval, the hot vertices form three convex polygons in three parallel planes. The middle polygon is a regular m -gon in the symmetry plane of the hyperboloid, and the other two are reflections of each other across that plane and are inscribed in the two hot circles. The Delaunay complex Δ_2 of the hot points is again degenerate, consisting of the above mentioned three polygons, which form the top and bottom facets of two drum-like polyhedra. The two drums are surrounded by a ring of four-sided pyramids alternating with tetrahedra.

Persistence Lemma B. *At any time in $(0, 2H^2h^2]$, the intersection of Δ_2 and D consists of the two rings of triangles forming the cylinder-with-a-waist and of their edges and vertices.*

Proof. The edges, polygons, and polyhedra in Δ_2 that do not belong to the cylinder-with-a-waist have their dual Voronoi polygons, edges, and vertices either in the symmetry plane or the symmetry axis of the hyperboloid. For the reason mentioned in the proof of the Persistence Lemma A, these polygons, edges, and vertices are contained in the union of balls bounded by sandwiching spheres of points of F inside the hot sphere. The corresponding edges, polygons, and polyhedra of Δ_2 thus do not belong to D .

The remainder of the proof establishes that the triangles uvw and vwx belong to D . Immediately after time $t = 0$ this is true because the triangles in the double-cup belonged to D immediately before time $t = 0$. At times $0 < t \leq t_1$ this is true because any violation is prevented by the algorithm before it occurs. $\square$

The two Persistence Lemmas also hold for the reverse metamorphosis, which changes a one-sheeted into a two-sheeted hyperboloid. To see this, run time backwards and exchange the arguments that establish that the two special configurations are subcomplexes of D when they are first created. These arguments are contained in the respective last paragraphs of the two proofs.

Summary. The two Persistence Lemmas establish 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.

11. Discussion

This paper describes a dynamic algorithm for maintaining the triangulation of a deforming skin surface by adapting it to changing shape, curvature, and topology.

Abstract Interface. The algorithm uses detailed knowledge about the skin surface to avoid pitfalls, such as insufficient quality of approximation, small angles, and wrong connections. In an effort to understand the extent to which the algorithm can be generalized, we may ask how much knowledge about the surface the algorithm really needs. Can we list axioms for a deforming surface that imply the applicability of the algorithm? To make this a worthwhile exercise, one would of course hope that the class of surfaces and deformations defined by the axioms is significantly larger than the class of skin surfaces and the growth model.

A key idea of the algorithm is to keep the density of vertices roughly proportional to one over the maximum curvature. This is only possible if that measure of curvature satisfies a one-sided Lipschitz condition, like that stated in the Curvature Variation Lemma. It is conceivable that the sampling density can be based on other expressions of curvature that satisfy a one-sided Lipschitz condition. The local feature size proposed in [3] is a candidate for such an expression, but it is usually not easy to compute. Another important cornerstone of the algorithm is the predictability of metamorphoses. Without knowing when and where a metamorphosis will happen, we would have to resort to relatively expensive collision detection and resolution algorithms.

Implementation. Where do we go from here? The first job on the agenda is the rigorous implementation of the dynamic skin triangulation algorithm. The first author of this paper

has already taken steps in that direction, partially reusing prior software on alpha shapes [9] and on computing Betti numbers [6]. It will be interesting to study the algorithm experimentally and measure the influence of design decisions on its performance. For example, the bounds on the constants C , Q controlling the curvature adaptation algorithm derived in this paper are all conservative. Perhaps it is possible to relax the requirements a fair amount without compromising the correctness of the algorithm.

It would be interesting to modify the algorithm to sample points with local density roughly proportional to $\sqrt{1/\kappa}$ rather than to $1/\kappa$. A density like that would lead to a uniformly approximating triangulation. In other words, the distance between the triangulation and the skin surface would be bounded by a constant independent of curvature. As another bonus, the vertices would be distributed without accumulation points at singularities with infinite curvature. As a corresponding drawback, the triangulation would suffer from violations of the closed ball property, which is crucial to the algorithm as described in this paper.

Deformation. Recall that the skin surface is defined by a finite set of spheres in $\mathbb{R}^3$. It is deformed by manipulating this set. The simplest type of deformation is plain growth, which is generated by increasing all radii continuously and simultaneously in a way that preserves the mixed complex. The next more complicated type of deformation is the one described in [4]. It can be used to morph one skin surface into any other skin surface by a combination of moving and growing/shrinking of spheres. This is done in a way that preserves the combinatorics but not the geometric realization of the mixed complex. Metamorphoses in the morphing model are only slightly more difficult to predict than in the growth model. The maintenance of the skin triangulation is however made more complicated by the occurrence of new types of degenerate metamorphoses, such as the ones that progress to the critical point and are then reversed. In any physical simulation, where the deformation depends on forces unrelated to the combinatorics of the mixed complex, we will also have to maintain that complex dynamically. The computational overhead is negligible if we use a dynamic algorithm for maintaining the Delaunay complex [11]. Note that a single time step in the simulation may jump over any number of metamorphoses and other changes. A reasonable approach to bridging the gaps in time is to connect any two contiguous time slices with a deformation in the morphing model mentioned above.

An important question in this context is the inverse problem. How do we construct the skin surface that best approximates a given shape, and how do we maintain such an approximation? We need modeling operations that allow us to increase or decrease the complexity of the surface locally. The former can be done by doubling spheres and continuously moving them apart. The latter can be done by the inverse operation, which moves spheres together and merges them into one.

Acknowledgments

The authors of this paper thank Robert Haber and the Center for Process Simulation and Design at the University of Illinois for stimulating discussions and for providing the practical background motivating the design of a dynamic meshing algorithm for

deforming surfaces. They also thank two anonymous referees for their constructive criticism.

Appendix

Figure 20 gives an overview of the proof structure used in this paper. Nodes represent properties, lemmas, invariants, and conditions. Arcs represent logical dependencies directed from top to bottom. All abbreviations are explained below. Table 2 provides a list of notation used in this paper.

(I)	Condition on ε	Section 4
(II)	Condition on C and Q	Section 8
(III)	Condition on C and Q	Section 8
(IV)	Condition on h, ℓ, m	Section 10
(V)	Condition on h, ℓ, m	Section 10
[L]	Lower Size Bound	Section 6
[U]	Upper Size Bound	Section 6
CcL	Circumcenter Lemma	Section 8
CSL	Curvature Sandwich Lemma	Section 3
CVL	Curvature Variation Lemma	Section 3
ENL	Edge Normal Lemma	Section 4
GHT	General Homeomorphism Theorem	Section 4
HSL	Hot Spot Lemma	Section 7
IcL	Iso-curvature Lemma	Section 3
LDC	Long Distance Claim	Section 4
MAL	Minimum Angle Lemma	Section 6
NSL	No-Short-Edge Lemma	Section 8

Fig. 20. Logical structure of the proof that the algorithm in this paper constructs a triangulation of the skin surface.

Table 2. Notation for important geometric concepts, functions, variables, and constants.

x_1, x_2, x_3	Coordinates
$\pi_{\hat{a}}: \mathbb{R}^3 \rightarrow \mathbb{R}$	Weighted (square) distance function
$\hat{a} = (a, A)$	Zero-set of $\pi_{\hat{a}}$, sphere with center a and radius A
$\alpha\hat{a} + \beta\hat{b}$	Zero-set of $\alpha\pi_{\hat{a}} + \beta\pi_{\hat{b}}$
$\sqrt{\hat{a}} = (a, A/\sqrt{2})$	Zero-set of $\pi_{\hat{a}} - \pi_{\hat{a}}(a)/2$
$F = \text{skin } \mathcal{A}$	Skin surface, $\text{env } \sqrt{\text{conv } \mathcal{A}}$
$\text{body } \mathcal{A}$	Body bounded by skin $\mathcal{A}$
$v_{\mathcal{X}}$	Intersection of Voronoi polyhedra
$\delta_{\mathcal{X}}$	Delaunay simplex
k	Dimension of Delaunay simplex
$\mu_{\mathcal{X}}$	Mixed cell, $(v_{\mathcal{X}} + \delta_{\mathcal{X}})/2$
$z_{\mathcal{X}}$	Center, $\text{aff } v_{\mathcal{X}} \cap \text{aff } \delta_{\mathcal{X}}$
$\kappa: F \rightarrow \mathbb{R}$	Maximum curvature function
$\varrho: F \rightarrow \mathbb{R}$	Length scale function, $1/\kappa$
$\mathbf{n}_x, \mathbf{n}_{xyz}$	Normal vector at x , of xyz
$\mathbf{t}_{xy}$	Tangent vector, $(y - x)/\ y - x\ $
S_x, T_x	Sandwiching spheres at $x \in F$
R_{ab}, R_{abc}	Radius or size of edge, triangle
ϱ_{ab}	Target size, $\max\{\varrho(a), \varrho(b)\}$
ϱ_{abc}	Target size, $\min\{\varrho(a), \varrho(b), \varrho(c)\}$
ε	Constant controlling sampling density
Q	Constant controlling triangle quality
C	Constant controlling approximation
$M: \mathbb{R}^3 \rightarrow \mathbb{R}$	Trajectory of skin surface
φ_{ij}	Diffeomorphism between skins
V	ε -Sampling of F
D	Restricted Delaunay triangulation
Δ	Delaunay complex
H	Radius of hot ball
$\mathbb{R}_H^3$	Hot portion of space
F_H	Hot portion of skin surface, $F \cap \mathbb{R}_H^3$
h	Constant triggering a metamorphosis
ℓ, m	Constant numbers of hot vertices
t_0, t_1	Start/end time of a metamorphosis
R_0, R_1	Start/end radius of hot circles

NVL	Normal Variation Lemma	Section 3
PLA	Persistence Lemma A	Section 10
PLB	Persistence Lemma B	Section 10
SDC	Short Distance Claim	Section 4
SgL	Sampling Lemma	Section 8
SHT	Special Homeomorphism Theorem	Section 10
SwP	Sandwich Property	Section 2
TNL	Triangle Normal Lemma	Section 4
TsL	Transition Lemma	Section 10
VEL	Voronoi Edge Lemma	Section 4
VGL	Voronoi Polygon Lemma	Section 4
VHL	Voronoi Polyhedron Lemma	Section 4

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. P. S. Alexandrov. *Combinatorial Topology*. Dover, New York, 1956.
3. N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.* **22** (1999), 481–504.
4. H.-L. Cheng, H. Edelsbrunner, and P. Fu. Shape space from deformation. In *Proc. 6th Pacific Conf. Comput. Graphics Appl.*, Singapore, 1998, pp. 104–113.
5. T. E. Creighton. *Proteins. Structures and Molecular Principles*. Freeman, New York, 1984.
6. 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.
7. H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.
8. H. Edelsbrunner. Deformable smooth surface design. *Discrete Comput. Geom.* **21** (1999), 87–115.
9. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graph.* **13** (1994), 43–72.
10. H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.* **7** (1997), 365–378.
11. M. A. Facello. Geometric techniques for molecular shape analysis. Ph.D. thesis, Report UIUCDCS-R-96-1967, Dept. Comput. Sci., Univ. Illinois, Urbana, Illinois, 1996.
12. A. Gray. *Modern Differential Geometry of Curves and Surfaces*. CRC Press, Boca Raton, Florida, 1993.
13. L. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graph.* **4** (1985), 74–123.
14. J. Milnor. *Morse Theory*. Princeton University Press, Princeton, New Jersey, 1963.
15. T. E. Morthland, P. E. Byrne, D. A. Tortorelli, and J. A. Dantzig. Optimal riser design for metal castings. *Metall. Material Trans.* **26B** (1995), 871–885.
16. D. Pedoe. *Geometry: a Comprehensive Course*. Dover, New York, 1988.
17. N. Provatas, N. Goldenfeld, and J. A. Dantzig. Adaptive grid methods in solidification microstructure modeling. *Phys. Rev. Lett.* **80** (1998), 3308–3311.

Received July 19, 2000, and in revised form January 6, 2001. Online publication March 26, 2001.