The Medial Axis of a Union of Balls

Nina Amenta [a,1] Ravi Krishna Kolluri [a,1]

[a] *University of Texas, Computer Sciences Dept. , Austin, TX 78712, USA*

Abstract

We present an algorithm for computing the exact interior medial axis of a union of balls in $\mathbb{R}^d$. Our algorithm combines the simple characterization of this medial axis given by Attali and Montanvert with the combinatorial information provided by Edelsbrunner's α-shape. This leads to a simple algorithm, which we have implemented for $d = 3$.

1 Introduction

There is considerable interest, in mesh generation, computer graphics, computer vision and medical imaging, in describing the shape of an object by its medial axis. But computing the medial axis of a three-dimensional object is difficult in general; it can be done for polyhedra using exact arithmetic and computational algebra [4].

Attali and Montanvert [1] showed that, in contrast, the medial axis of a three-dimensional union of balls has a simple structure. This is an important special case, since *any* three-dimensional shape can be approximated by a finite union of balls; we recently showed [5] that given a certain well-defined sample of points P on an object surface, the union of a subset of the Voronoi balls of P is a good approximation of the original object, for which the errors in position and normal on the surface are everywhere bounded. Thus, given a good sample from an object surface, we can compute the exact medial axis of a nearby union of balls. In the quite common case in which the original object is known only by the sample points, this is as good an approximation of the medial axis as we are likely to get.

The elegant characterization of the medial axis of a union of balls given by Attali and Montanvert does not lead immediately to a robust algorithm, however. Using the beautiful duality relationship between a union of balls and

[1] Both authors are supported by NSF/CCR-9731977, and the first author is supported by an Alfred P. Sloan Foundation Fellowship.

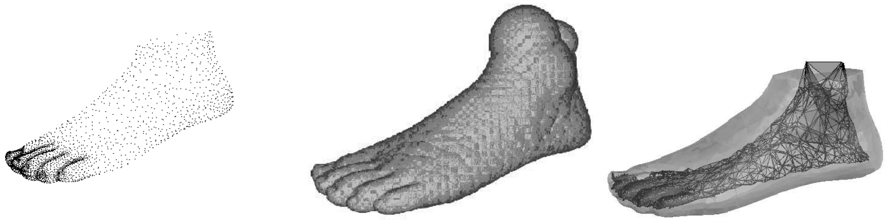

Fig. 1. Left, set of sample points on an object surface; Middle, a union of balls which accurately represents the object; Right, the medial axis of the union of balls shown inside the object surface.

its *α-shape*, developed by Edelsbrunner [3], we give a more combinatorial and algorithmic version of this characterization. This gives us an algorithm for the construction of the medial axis of the union of balls, which we have implemented in $\mathbb{R}^3$, using the CGAL library.

A numerically robust exact-arithmetic implementation of our construction would involve orientation tests of quite high degree: we need to compute the Delaunay triangulation of a set of points, each defined as an intersection of three of the input spheres. Instead, we compute the coordinates of these points as accurately as we can, represent them as floating-point numbers, and then use them as input to the "off the shelf" robust three-dimensional Delaunay triangulation program in CGAL. This seems to work well on the examples we have tried.

2 Definitions

Let $\mathcal{B}$ be a set of d-dimensional balls in $\mathbb{R}^d$, and let $\mathcal{U}$ be their union. We assume that the balls are in general position. This simplifies the following:

Definition: A *k-face* of $\mathcal{U}$ is a subset of the boundary $\delta\mathcal{U}$ contained in the intersection of a subset of $\mathcal{B}$ of cardinality exactly $(d-k)$.

Note that a k-face is a subset of a k-sphere and need not be connected.

Definition: The *vertices* of $\mathcal{U}$ are the points forming the 0-faces of $\mathcal{U}$. Let $\mathcal{V}$ be the set of vertices of $\mathcal{U}$.

A union of balls and its set of vertices are shown in Figure 2.

Edelsbrunner [3] studied the dual shape of $\mathcal{U}$, the *α-shape* $\mathcal{S}$. The *power diagram Pow*($\mathcal{B}$) of a set $\mathcal{B}$ of balls is a kind of Voronoi diagram, dividing space into polyhedral cells, each of which represents the points "closest" to a particular input ball $b \in \mathcal{B}$ (see [3] for formal definitions and many properties). Like the usual Voronoi diagram, *Pow*($\mathcal{B}$) has a dual Delaunay triangulation, know as a *regular triangulation*, consisting of the simplices connecting the centers of every set T of balls forming a face of *Pow*($\mathcal{B}$). Restricting the sub-

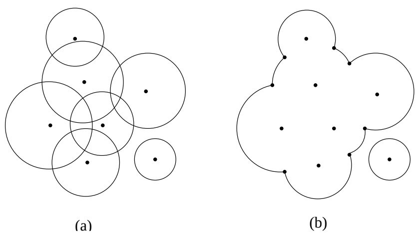

Fig. 2. A set of balls $\mathcal{B}$ in $\mathbb{R}^2$ and the union $\mathcal{U}$ of the balls in $\mathcal{B}$. The vertices of $\mathcal{U}$ are shown as black points on the boundary of the union.

division of space defined by $Pow(\mathcal{B})$ to the union $\mathcal{U}$ gives a subdivision of $\mathcal{U}$ (see Figure 4a) into cells. The subset of the faces of $Pow(\mathcal{B})$ participating in the subdivision of $\mathcal{U}$ corresponds to a subset of the simplicies forming the dual regular triangulation; this subset of simplicies forms the *α-complex* of $\mathcal{B}$. See Figure 4b. The underlying space of this complex (ie. the union of all its simplices) is the *α-shape* $\mathcal{S}$; see Figure 4c.

Edelsbrunner showed that the structure of the α-shape reflects the combinatorial structure of the boundary $\delta\mathcal{U}$ of the union. Let s_T be any face on the boundary $\delta\mathcal{U}$, formed by the intersection of the boundaries of k balls in a subset $T \subseteq \mathcal{B}$; s_T has dimension $d-k$. Then the k-dimensional simplex σ_T is part of the boundary $\delta\mathcal{S}$ of the α-shape. Consult Figure 4c for an example. We say that the face s_T of $\delta\mathcal{U}$ and the face σ_T of $\mathcal{S}$ are *duals*.

Later, we will need to distinguish the full-dimensional parts of the α-shape from the rest.

Definition: A face in the boundary $\delta\mathcal{S}$ of the α-shape is *singular* if it is not a face of a a d-dimensional simplex in the α-complex $\mathcal{K}$. Consider removing the singular faces from the α-shape $\mathcal{S}$; the connected components of the remaining set are the *regular components*, the full-dimensional parts of the α-shape.

A regular component $\mathcal{C}$ of the α-shape $\mathcal{S}$ is a solid, but it need not be a regular solid; that is, it need not have a boundary which is a piecewise-linear manifold, and hence, for instance, a $(d-2)$-face might be contained in more than two $(d-1)$-faces. Figure 3 shows a three-dimensional example.

Definition: The *interior medial axis* (or simply the *medial axis*) $\mathcal{X}$ of an object O is the closure of the set of points $m \in O$ such that m has at least two closest points on the surface δO.

Consider a ball b, contained in O and touching δO in at least two points. The definition of the medial axis implies that the center p of b is a point on the medial axis. We can now state the elegant characterization of the medial axis of the union of balls given by Attali and Montanvert [1].

Theorem 1 *Let $\mathcal{U}$ be a union of balls, let $\mathcal{V}$ be its vertices and let $\mathcal{S}$ be its α-shape. The medial axis $\mathcal{X}$ of $\mathcal{U}$ consists of*
(i) the singular faces of $\mathcal{S}$, and

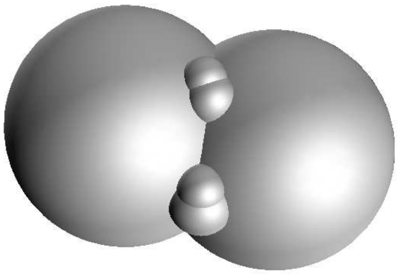

Fig. 3. The α-shape of this union of balls consists of a single regular component $\mathcal{C}$, made up of two tetrahedra sharing the single long edge between the centers of the two big balls.

(ii) the subset of the Voronoi diagram $Vor(\mathcal{V})$ *whose closest point on the boundary* $\delta\mathcal{U}$ *is a vertex in* $\mathcal{V}$.

The difficulty with using this theorem directly in the computation of the medial axis $\mathcal{X}$ of $\mathcal{U}$ is that in dimensions three and higher, unlike $\mathbb{R}^2$, the Voronoi diagram of $\mathcal{V}$ can be quite complicated, including Voronoi vertices interior to $\mathcal{U}$ but not belonging to $\mathcal{X}$, and higher-dimensional faces of $Vor(\mathcal{V})$ which partially belong to $\mathcal{X}$. Testing condition (i), above, is easy, but it is not obvious how condition (ii) should be tested. Computing the location of vertices, edges and faces of $Vor(\mathcal{V})$, and determining which of them, or which parts of them, are closer to a vertex in $\mathcal{V}$ than to any other point on $\delta\mathcal{U}$, is fraught with numerical difficulties. This motivates our more combinatorial characterization of the subset of $Vor(\mathcal{V})$ belonging to the medial axis, which is given in the next section.

We should also note that the point set $\mathcal{V}$ is quite degenerate; it generally will contain many sets of $k > (d+1)$ points which lie on the same ball, many sets of $k > d$ points lying on the intersection of two balls, etc. Handling this rather degenerate structure combinatorially, rather than numerically, seems essential to producing a robust algorithm.

3 Geometric theorems

The following theorem gives our characterization of the medial axis of $\mathcal{U}$. Recall that the regular components are the connected components of the full-dimensional parts of the α-shape.

Theorem 2 *Let* $\mathcal{U}$ *be a union* $\mathcal{U}$ *of balls in* $\mathbb{R}^d$, *let* $\mathcal{V}$ *be the vertices of* $\delta\mathcal{U}$ *and let* $\mathcal{S}$ *be the* α*-shape of* $\mathcal{U}$. *The medial axis* $\mathcal{X}$ *of* $\mathcal{U}$ *consists of*
(i) the singular faces of $\mathcal{S}$, *and*
(ii) the subset of $Vor(\mathcal{V})$ *which intersects the regular components of* $\mathcal{S}$.

Our goal in this section is the proof of Theorem 2. We begin with some observations about the properties of the α-shape, which follow from the proof of

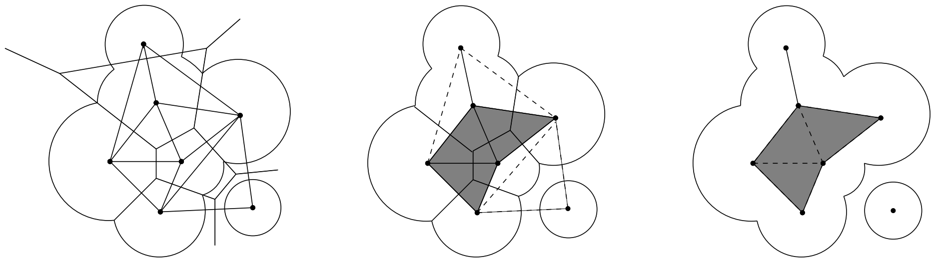

Fig. 4-a. Fig. 4-b. Fig. 4-c.

Fig. 4. Left, the regular triangulation and the power diagram of the union of balls shown in Figure 2. Middle, the α-*complex* $\mathcal{K}$ of a union $\mathcal{U}$ of balls. Note that the α-complex is a subset of the regular triangulation. Right, the α-*shape* $\mathcal{S}$ which is the underlying space of $\mathcal{K}$. The line segment and the center of the isolated ball are singular components of $\mathcal{S}$, and the remainder is a regular component.

the homotopy equivalence of $\mathcal{S}$ and $\mathcal{U}$ given by Edelsbrunner [3].

First, we note that points in $\mathcal{U} - \mathcal{S}$ cannot belong to the medial axis of $\mathcal{U}$.

Observation 3 *For each point $x \in \mathcal{U} - \mathcal{S}$, there is a unique point $u \in \delta\mathcal{U}$ such that for all $v \in \delta\mathcal{U}$, $d(u, x) < d(v, x)$.*

See point u in Figure 5.

Also, as observed by Edelsbrunner [3], points in a simplex σ_T on the boundary $\delta\mathcal{S}$ of the α-shape are closest to the dual face s_T on $\delta\mathcal{U}$:

Observation 4 *If $y \in \sigma_T$, where σ_T is a simplex in $\delta\mathcal{S}$, then y is the center of a d-ball b_y with the dual face $s_T \subseteq \delta b_y \subseteq \mathcal{U}$.*

See points v and w in Figure 5.

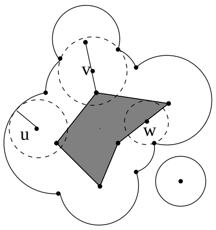

Fig. 5. Point u in $\mathcal{U} - \mathcal{S}$ has a unique closest point in $\delta\mathcal{U}$. Points v and w lie on simplices in $\delta\mathcal{S}$ and have their closest points in $\delta\mathcal{U}$ on the corresponding dual faces.

Observation 5 *A singular $(d-1)$-face in $\mathcal{S}$ is dual to a 0-face containing two vertices of $\mathcal{U}$. A $(d-1)$-face that belongs to a regular component of $\mathcal{S}$ is dual to a single vertex of $\mathcal{U}$.*

Now we will show that certain parts of the α-shape have to belong to the medial axis.

Lemma 6 *Any k-simplex of $\delta\mathcal{S}$, for $(d-1) > k \geq 0$ belongs to the medial axis $\mathcal{X}$ of $\mathcal{U}$.*

Proof: Consider a k-face σ_T of $\mathcal{S}$, with $k < (d-1)$. The dimension of s_T, the dual face of σ_T, is greater than zero. By Observation 4, any point x contained in σ_T is the center of a ball b with $s_T \subseteq \delta b \subseteq \mathcal{U}$. Since the dimension of s_T is at least one, δb contains more than one point of $\delta\mathcal{U}$, and x is a point of the medial axis.
□

The following is due to Attali and Montanvert; we include it for completness.

Lemma 7 *Any point belonging to a singular face of $\mathcal{S}$ belongs to the medial axis $\mathcal{X}$ of $\mathcal{U}$.*

Proof: By Lemma 6, any singular face σ_T of dimension $k < (d-1)$ belongs to the medial axis. So let us consider the remaining case, a singular simplex σ_T of dimension $(d-1)$. Let x be a point contained in σ_T. By Observation 5, the dual face s_T consists of two vertices $p, q \in \mathcal{V}$. Again by Observation 4, p, q are closest points to x in $\delta\mathcal{U}$, and x belongs to the medial axis.
□

The Lemma 6 showed that, for $k < (d-1)$, every point on a k-face on the boundary of a regular component $\mathcal{C}$ belongs to the medial axis. We now consider the interior of $\mathcal{C}$ and the interiors of the $(d-1)$-faces on its boundary. We begin by showing that for any point in $\mathcal{C}$, some closest point in $\delta\mathcal{U}$ is a vertex of $\mathcal{U}$.

Lemma 8 *Let $\mathcal{C}$ be a regular component of $\mathcal{S}$. Then for any point $x \in \mathcal{C}$, either*
a) x has exactly one closest point p on $\delta\mathcal{U}$, and p is a vertex of $\mathcal{U}$, or
b) x has more than one closest point on $\delta\mathcal{U}$, at least two of which are vertices of $\mathcal{U}$.

Proof: We first consider any point x in the interior of $\mathcal{C}$ (eg. point v in Figure 6). We claim that any closest point to x must be a vertex of $\delta\mathcal{U}$. Assume for the purpose of contradiction that there is a point p closest to x such that $p \in s_T$ for some face s_T of dimension greater than zero. Consider the ball b centered at x and touching p. Since p is a closest point to x on $\delta\mathcal{U}$, $b \subseteq \mathcal{U}$. The segment px intersects the boundary of $\mathcal{S}$ in some point y. By Observation 4, there is a ball b_y centered at y with $s_T \subseteq \delta b_y \subseteq \mathcal{U}$, and since $y \in px$, point p is also a closest surface point to y. Hence y must belong to the face σ_T of $\delta\mathcal{S}$ dual to s_T, and every point of σ_T is a closest surface point to y. Since y is closer to p than x is, $b_y - p$ lies in the interior of $b - p$, so that the set $s_T - p$, of dimension greater than zero, lies in the interior of $b - p$. Since s_T is part of $\delta\mathcal{U}$, this contradicts $b \subseteq \mathcal{U}$. Thus, while a point x in the interior may have one closest point on $\delta\mathcal{U}$ (a), or more than one (b), in either case every closest point must be a vertex of $\delta\mathcal{U}$.

Now we consider the case in which $x \in \delta\mathcal{C}$. Let σ_T be the simplex of smallest dimension in $\delta\mathcal{C}$ containing x. If σ_T is a $(d-1)$-simplex, then the only closest point to x in $\delta\mathcal{C}$ is the dual face s_T, which is a vertex of $\delta\mathcal{U}$, and we are done.

Otherwise, σ_T is incident to at least two $(d-1)$-simplices in $\delta\mathcal{C}$, so that its dual face s_T must contain at least two vertices of $\delta\mathcal{U}$, both of which are closest points to x. This is case (b); note that x might have other closest points in s_T which are not vertices (eg. point u in Figure 6).
□

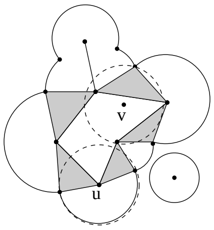

Fig. 6. The shaded regions connect each face in the regular component to its dual face in $\mathcal{U}$. Point v has a unique closest vertex in $\delta\mathcal{U}$. For u, all the points lying on the dotted arc, including the two vertices at the end points, are the closest points in $\delta\mathcal{U}$.

We are now in a position to prove our theorem characterizing the medial axis $\mathcal{X}$ of $\mathcal{U}$.

Proof of Theorem 2: First, we show that any point which either is a singular face of $\mathcal{S}$ or belongs to both a regular component of $\mathcal{S}$ and to $Vor(\mathcal{V})$ must be a point of the medial axis. By Lemma 7, any point in a singular face of $\mathcal{S}$ belongs to the medial axis. Any point x in a regular component belonging to $Vor(\mathcal{V})$ clearly has more than one closest vertex on $\delta\mathcal{U}$. Lemma 8 implies that no other point on $\delta\mathcal{U}$ is closer to x, and hence x must belong to the medial axis of $\mathcal{U}$.

Now consider a point m in the medial axis of $\mathcal{U}$; it cannot be a point in $\mathcal{U}-\mathcal{S}$, by Observation 3. If m belongs to a singular face we are done. Otherwise m belongs to some regular component, and it must have more than one closest point in $\mathcal{V}$, by Lemma 8, so that it belongs to $Vor(\mathcal{V})$.
□

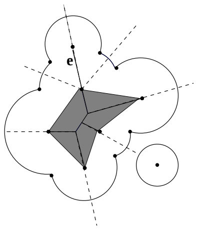

Fig. 7. The medial axis of the union of balls consists of the singular faces - the edge e and the center of the isolated ball - and the the edges of $Vor(\mathcal{V})$ inside the (shaded) regular component, shown as solid lines.

4 Combinatorial Tests

This version of the characterization of the medial axis leads to a (mostly) combinatorial algorithm for its computation. The idea of the algorithm is that we can find the portion of $Vor(\mathcal{V})$ contained inside each regular component C by tracing the edge-skeleton of $Vor(\mathcal{V})$, and computing its intersection with the boundary δC.

We want to use combinatorial information to compare $Vor(\mathcal{V})$ with $\mathcal{C}$. Each simplex σ_T in $\delta\mathcal{C}$ is defined as the convex hull of the centers of a set T of balls; we associate T of course with σ_T. Each vertex v_T in $\mathcal{V}$ lies on the surface of a set T of d balls; we similarly associate T with v_T. Finally, each edge e_V in $Vor(\mathcal{V})$ is associated with a set of $V \subseteq \mathcal{V}$ of $k \geq d$ vertices on $\delta\mathcal{U}$, and hence with the sets of balls determining each of those vertices.

To develop our combinatorial tests, we consider the ways in which simplices of $\delta\mathcal{C}$ interact with $Vor(\mathcal{V})$. We show in this section that an edge e_V of $Vor(\mathcal{V})$ can only intersect $\delta\mathcal{C}$ in two specific ways, which are easy to detect by examining the sets of input balls associated with e_V.

First, we establish that the $(d-1)$-faces of the boundary $\delta\mathcal{C}$ of a regular component interact in a simple way with $Vor(\mathcal{V})$: each $(d-1)$ face divides a Voronoi cell neatly in half.

Lemma 9 *The interior of any $(d-1)$-simplex $\sigma \in \delta\mathcal{C}$ is in the interior of the Voronoi cell of its dual vertex v.*

Proof: ¿From Observation 4 we know that for each point p in the interior of σ the unique closest vertex in $\mathcal{V}$ is v, the dual face of σ. Hence p belongs to the interior of the Voronoi cell of v.
□

Lemma 10 *Each k-simplex $\sigma_T \in \delta\mathcal{C}$ for $k < (d-1)$ is part of a face of $Vor(\mathcal{V})$ of dimension $\leq (d-1)$.*

Proof: Let $V \subset \mathcal{V}$ be the set of vertices that are dual to the $(d-1)$-simplices incident on σ_T. We have $|V| > 1$. Every $v \in V$ lies on the $(d-k)$ sphere dual to σ_T and containing s_T. ¿From Observation 4, for any point $y \in \sigma_T$, each $v \in V$ is a closest vertex in $\delta\mathcal{U}$, and hence y belongs to the Voronoi face f_V induced by the vertices in V. Since $|V| > 1$, face f_V has dimension at most $d-1$.
□

Theorem 11 *Let σ be a $(d-1)$-simplex on the boundary of $\delta\mathcal{C}$, and let $v \in \mathcal{V}$ be its dual vertex. Face σ divides the Voronoi cell $Vor(v)$ of v into two parts, one inside $\mathcal{C}$ and one outside.*

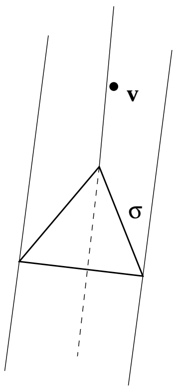

Fig. 8. Simplex σ partitions the Voronoi cell of v.

Proof: ¿From Lemma 10 each $(d-2)$-simplex in the boundary of σ belongs to the boundary of $Vor(v)$, and from Lemma 9 the interior of σ belongs to the interior of $Vor(v)$. The lemma follows because $Vor(v)$ is a topological ball.
□

As the following theorem shows, this simple structure implies that the intersection of edges of $Vor(\mathcal{V})$ with $\delta\mathcal{C}$ fall into two categories.

Lemma 12 *Let e be an edge of $Vor(\mathcal{V})$ intersecting $\delta\mathcal{C}$, and let σ be the face of $\delta\mathcal{C}$ of smallest dimension containing the intersection of e and $\delta\mathcal{C}$. Either*
a) σ is a vertex of $\mathcal{C}$, or
b) σ is identical to e.

Proof: Consider a $(d-1)$-face f of $\delta\mathcal{C}$ containing σ. By Theorem 11, f divides the Voronoi cell $Vor(v)$ of the vertex $v \in \mathcal{V}$ dual to f into into two parts, inner and outer. The faces of f similarly are formed by the intersection of f with the boundary of $Vor(v)$.

Since e is an edge of $Vor(\mathcal{V})$, e can intersect the cell $Vor(v)$ in one of two ways: either e intersects $Vor(v)$ in a single vertex (in which case we are done; σ must consist of that vertex), or e is an edge of $Vor(v)$. In the latter case, either f completely contains that edge (in which case we are done; σ and e are identical), or f intersects that edge in a vertex of $\delta\mathcal{C}$ (and we are done, with σ as that vertex).
□

When e intersects $\delta\mathcal{C}$ in a vertex we call e a *crossing edge*, while when e is itself an edge of $\delta\mathcal{C}$ we call e an *α-shape edge*. Figure 3 provides an example of a union of balls inducing an an α-shape edge: the long edge connecting the centers of the two large balls in the α-shape is also an edge of $Vor(\mathcal{V})$, equidistant from the four vertices on the intersection of the two large balls.

The following two lemmas show how crossing edges and α-shape edges can be identified combinatorially, by examining the sets of balls inducing features of

$Vor(\mathcal{V})$ and of $\mathcal{C}$.

Lemma 13 *Let e_V be an edge of $Vor(\mathcal{V})$ induced by the set $V \subseteq \mathcal{V}$. Let σ be a vertex of $\delta\mathcal{C}$, which is the center of input ball b. Edge e_V passes through σ if and only if the vertices in V all lie on the boundary of b.*

Proof: When all the vertices in V lie on ball b, edge e_V passes through σ, since every vertex in V is a closest vertex in $\mathcal{V}$ to σ . This is shown in Figure 9. The converse is true because σ is incident on at least d simplices of dimension $(d-1)$ in $\delta\mathcal{C}$, each of which is dual to a vertex in $\mathcal{V}$ on b, which is closest to σ. So any closest vertex to σ must also lie on b.
$\square$

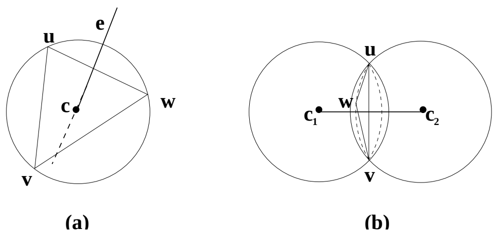

Fig. 9. The two ways in which a Voronoi edge interacts with $\delta\mathcal{C}$. Left, a Voronoi edge e_V passing through a single vertex of $\mathcal{C}$. The set vertices $V = \{u, v, w\}$ inducing the edge is also shown. Right, a Voronoi edge which is an edge of $\delta\mathcal{C}$ (this case does not occur in 2D; note that in 3D the set V of vertices inducing edge e_V would have to have cardinality at least four, as eg. in Figure 3).

Lemma 14 *Let e_V be an edge of $Vor(\mathcal{V})$ induced by the set $V \subseteq \mathcal{V}$. Let σ be an edge of $\delta\mathcal{C}$, which connects the centers of input balls b_1, b_2. Edge e_V is identical to σ if and only if the vertices in V lie on the intersection of $b_1 \cap b_2$ in $\delta\mathcal{U}$.*

Proof: If e_V is identical to σ then the vertices in V must be closest vertices to σ, and hence must lie on the face s of $\delta\mathcal{U}$ dual to σ (Observation 9), which is a subset of the intersection $b_1 \cap b_2$.

Conversely, assume all of the vertices in V lie on $b_1 \cap b_2$. Then the ball centers b_1, b_2 lie on Voronoi edge e_V, and $\sigma \subset e_V$. Now consider any $(d-1)$-face f of $\delta\mathcal{C}$ containing σ, dual to a vertex $v \in V$. Edge e_V is one of the edges of Voronoi cell $Vor(v)$. The face f divides $Vor(v)$ into two parts, and σ is an edge of f. So if σ contains part of Voronoi edge e_V then in fact σ contains all of e and e_V and σ are identical.
$\square$

We use Lemmas 13 and 14 to identify the crossing edges and the α-shape edges. For each e_V edge in $Vor(\mathcal{V})$, we compute the common intersection of the family of sets:

$$B = \bigcap \{T \mid v_T \in e_V\}$$

That is, we find the set of balls such that all of the vertices inducing e_V lie on all of the balls. If B consists of a single ball b, e_V is a crossing edge. If B contains a pair of balls b_1, b_2, e_V is an α-shape edge. If e_V does not cross the boundary $\delta\mathcal{C}$, then B will be empty.

5 Algorithm

We now summarize our algorithm for extracting the medial axis of a union of balls in $\mathbb{R}^3$. The main idea is to label every vertex of $Vor(\mathcal{V})$ as *belonging* to or *exterior* to the α-shape $\mathcal{S}$. Vertices on the boundary of $\mathcal{S}$ are counted as belonging to $\mathcal{S}$. A face of $Vor(\mathcal{V})$ for which all of the vertices belong to $\mathcal{C}$ can then be output as part of the medial axis. We label the vertices of $Vor(\mathcal{V})$ by traversing its edge-skeleton. We change the label we assign to vertices every time we determine that we have crossed the boundary $\delta\mathcal{C}$. We detect these crossings using the combinatorial test defined in the previous section.

In: A set $\mathcal{B}$ of balls. $\mathcal{S}$.
Out: The medial axis $\mathcal{X}$ of their union $\mathcal{U}$.

0: Compute the α-shape $\mathcal{S}$ of $\mathcal{B}$.
1: Compute the set of vertices $\mathcal{V}$ of $\mathcal{U}$ as the duals of the triangles of $\delta\mathcal{S}$.
2: Add all edges (singular and regular) of $\mathcal{S}$ to $\mathcal{X}$.
3: Add all singular triangles in $\mathcal{S}$ to X.
4: Compute $Vor(\mathcal{V})$
5: Determine all crossing and α-shape edges by examining the sets of balls inducing the vertices of their dual triangles.
6: Split every crossing edge at the ball center that it crosses.
7: Label all Voronoi vertices which are centers of balls in $\mathcal{B}$ as belonging to $\mathcal{S}$.
8: Label a Voronoi vertex at infinity as exterior to $\mathcal{S}$.
9: Iteratively label each Voronoi cell $Vor(v)$ with at least one labeled vertex x as follows:
  (i) All vertices that can be reached without crossing the α-shape get the label of x (belonging or exterior).
  (ii) The remaining unlabeled vertices of $Vor(v)$ get the opposite label (belonging if x is exterior, or exterior if x is belonging).
10: Add to $\mathcal{X}$ all faces of the Voronoi diagram whose vertices all belong to $\mathcal{S}$.
11: Compute the locations of the vertices of $Vor(\mathcal{V})$ which belong to the medial axis, and output the vertices and faces of the medial axis.

Model	Balls	Vertices	Regular	Alpha	Delaunay	Label	Total
FOOT	2941	8266	41	13	296	12	362
MANNEQUIN	3412	9054	51	15	339	14	419
FIST	5493	14984	78	24	643	23	768

Table 1
The running times of the medial axis algorithm in seconds. The column *Regular* shows the time required to compute the regular triangulation of the input set of balls, *Alpha* shows the time required to select the α-shape from the regular triangulation, *Delaunay* shows the time required to compute the Delaunay triangulation of the vertices in $\mathcal{V}$, and *Label* is the time required for selecting the parts of the Voronoi diagram belonging to the medial axis.

6 Implementation

We implemented our medial axis algorithm in $\mathbb{R}^3$ using the computational geometry algorithms library CGAL, in particular the three dimensional regular triangulation and Delaunay triangulation modules. CGAL provides templates for using different number types; we used filtered arithmetic [7], which performs computations in floating point and reverts to exact arithmetic when the precision is insufficient to make correct combinatorial decisions. Filtered arithmetic is slower than floating point, but Delaunay triangulations cannot in general be computed using naive floating point.

The one place in which inexact arithmetic is used is in our algorithm is in the computation of the positions of the points in $\mathcal{V}$. These positions can be quite unstable, eg. a vertex formed where three large balls just touch. We compute the vertex positions as best we can and represent them as floating point numbers, which we then use as input to the computation of $Vor(\mathcal{V})$

Since the positions of the points in $\mathcal{V}$ are not exact, our version of $Vor(\mathcal{V})$ is *nearly*, rather than perfectly, degenerate. This means that what would be degenerate faces in $Vor(\mathcal{V})$, eg. vertices adjacent to more than d faces of dimension $(d-1)$, are broken up into several nearly-coincident non-degenerate faces. Because we perform only combinatorial tests on these faces, this does not complicate our algorithm at all.

We compute the α-shape using the algorithm explained in [2]. For the α-shape computation we need the ortho-center of each simplex in the regular triangulation, for which we used the exact predicates of Jonathan Shewchuk [6]. The ortho-centers of the triangles are also used in computing the vertices $\mathcal{V}$ of $\mathcal{U}$. These vertices are dual to triangles on the boundary $\delta\mathcal{C}$; they lie on a line normal to the plane of the triangle and passing through its ortho-center.

For each regular component $\mathcal{C}$, we built the Voronoi diagram of $\mathcal{V}_\mathcal{C}$, the subset of $\mathcal{V}$ dual to the $(d-1)$-dimensional faces of $\delta\mathcal{C}$, separately. We then computed the intersection of $Vor(\mathcal{V}_\mathcal{C})$ with $\mathcal{C}$. This is appropriate since each point inside

$\mathcal{C}$ is closest to a point in $\mathcal{V}_{\mathcal{C}}$:

$$(Vor(\mathcal{V}_{\mathcal{C}}) \cap \mathcal{C}) = (Vor(\mathcal{V}) \cap \mathcal{C})$$

We compute these Voronoi diagrams separately since, when possible, it is faster to do several smaller Voronoi diagram computations rather than one big one.

We used the unions of balls generated by the algorithm described in [5] to test our implementation. We ran our experiments on an SGI ONYX machine with a gigabyte of memory.

Table 5 breaks down the running time required by each step of the algorithm. It is clear that most of the time was spent in the computation of the Delaunay triangulation of $\mathcal{V}$. The percentage of time required for the α-shape computation and for labeling vertices is quite small, and the time required for the regular triangulation steps is considerably smaller than that required by the Delaunay triangulation.

The explanation for this is that $\mathcal{V}$ is highly degenerate. We believe that as a result the floating point filter was largely ineffective, and most of the computation was done using the exact arithmetic, which is much slower.

7 Discussion

Ideally we would like to restrict the numerical computations in our algorithm to those within the (trusted) Delaunay triangulation and regular triangulation routines, or to other exact arithmetic functions. But we have not succeeded in this, since we compute the positions of the vertices in $\mathcal{V}$ numerically. While it is possible that this introduces errors in the combinatorial structure of $Vor(\mathcal{V})$, none the less, the implementation seems to work well.

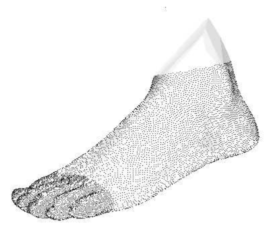

Fig. 10-a.

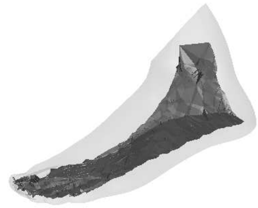

Fig. 10-b.

Fig. 10. Left, the sample points from which the FOOT model was generated. Right, the medial axis of the FOOT. Notice the line which joins the components corresponding to the top elongated portion and flat portions of the foot.

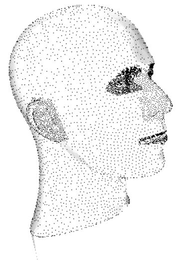

Fig. 11-a.

Fig. 11-b.

Fig. 11. Left, the sample points from which the MANNEQUIN model was generated. Right, the medial axis of the MANNEQUIN. Since the top region of the head is approximated by a few balls, the medial axis of this part is a small extension which is seen at the top.

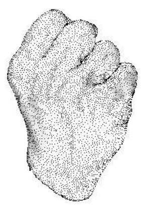

Fig. 12-a.

Fig. 12-b.

Fig. 12. Left, the sample used to generate the FIST model. Right, the medial axis of the FIST. Notice the ridges that correspond to the knuckles and the fingers in the fist.

References

[1] D. Attali and A. Montanvert. Computing and Simplifying 2D and 3D Continuous Skeletons, *Computer Vision and Image Understanding*, 67(3) (1997), pp. 261–273.

[2] H. Edelsbrunner. Weighted alpha shapes, *Tech. Rep. UIUCDCS-R-92-1760,*,Dept. of Computer Science, Univ. of Illinois at Urbana-Champaign, Urbana,Ill.

[3] H. Edelsbrunner. The union of balls and its dual shape, *Discrete Computational Geometry* 13(1995), 415-440.

[4] T. Culver, J. Keyser, and D. Manocha. Accurate computation of the medial axis of a polyhedron, UNC Technical Report TR98-038. Appeared in *Proc. ACM Solid Modeling*, pp,179-190,(1999).

[5] Nina Amenta and Ravi Kolluri. Accurate and efficient unions of balls. *Proceedings of the 16th Annual ACM Symposium on Computational Geometry*, pp, 119-128, (2000).

[6] Jonathan Richard Shewchuk. Adaptive Precision Floating-Point Arithmetic and Fast Robust Predicates for Computational Geometry. `http://www.cs.cmu.edu/~quake/robust.html`

[7] H. Brnnimann, C. Burnikel, and S. Pion. Interval arithmetic yields efficient arithmetic filters for computational geometry. *Proceedings of the 14th Annual ACM Symposium on Computational Geometry*, pp, 165-174, (1998).