Voronoi Diagrams on the Sphere*

Hyeon-Suk Na† Chung-Nim Lee† Otfried Cheong‡

June 22, 2001

Abstract

Given a set of compact sites on a sphere, we show that their spherical Voronoi diagram can be computed by computing two *planar* Voronoi diagrams of suitably transformed sites in the plane. We also show that a planar furthest-site Voronoi diagram can always be obtained as a portion of a nearest-site Voronoi diagram of a set of transformed sites. Two immediate applications are an $O(n \log n)$ algorithm for the spherical Voronoi diagram of a set of circular arcs on the sphere, and an $O(n \log n)$ algorithm for the furthest-site Voronoi diagram for a set of circular arcs in the plane.

1 Introduction

Voronoi diagrams belong to the computational geometer's favorite structures. They arise in nature and have applications in many fields of science [3, 15].

While Voronoi diagrams in the plane have been studied extensively, using different notions of sites and metrics, little is known for other geometric spaces. We are interested in the Voronoi diagram $\mathcal{SV}(\mathcal{U})$ of a set of sites $\mathcal{U}$ on a sphere $\mathcal{S}$, with the Euclidean distance on the surface of the sphere. This is a rather natural setting, considering that we are living on the surface of a large sphere. (Sugihara [15] gives a military motivation: the dominance regions of the various air bases in the world.) Note that for a fixed radius sphere $\mathcal{S}$, the geodesic distance of two points on $\mathcal{S}$ is a strictly monotone function of their Euclidean distance, and so the Voronoi diagram on $\mathcal{S}$ is in fact the intersection of the 3-dimensional Voronoi diagram of $\mathcal{U}$ with $\mathcal{S}$. Alternatively, we can use a stereographic projection to map $\mathcal{S}$ to a plane. The projection induces a metric on the plane, where the distance of two points is the geodesic distance of the corresponding points on $\mathcal{S}$. The Voronoi diagram of $\mathcal{U}$ can be obtained as a planar Voronoi diagram under this induced metric.

Unfortunately, these two observations do not help us to compute $\mathcal{SV}(\mathcal{U})$ efficiently. The 3-dimensional Voronoi diagram can have quadratic complexity even for point sets, and no algorithm appears to be known for the planar diagram under the metric induced by the stereographic projection. The difficulty is that the bisector of two sites under this metric is a closed curve. Very little is known about Voronoi diagrams under metrics with closed bisectors. The most general approach to the computation of planar Voronoi diagrams, the *abstract Voronoi diagram* by Klein [11, 12], requires that bisectors partition the plane into two unbounded regions. Algorithms that do work for sites with closed bisectors (such as the Voronoi diagram algorithm for curves by Alt and

*The authors wish to acknowledge the financial support of the Korea Research Foundation made in the program year of 1998, and the support of the Hong Kong Research Grant Council. Part of this research was done when O. Cheong was at Hong Kong University of Science & Technology, during a visit of H.-S. Na supported by the HKRGC.

†Department of Mathematics, Pohang University of Science and Technology, E-mail: hsnaa@postech.ac.kr, cnlee@postech.ac.kr

‡Institute of Information and Computing Sciences, Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, the Netherlands. E-mail: otfried@cs.uu.nl.

Schwarzkopf [1]) carefully work around the problem by cutting sites and adding point sites at the cuts.

However, computing the Voronoi diagram $\mathcal{SV}(\mathcal{U})$ of a set $\mathcal{U}$ of *points* on the sphere turns out to be surprisingly easy. One of the very first works in the then newly appearing field of computational geometry was Brown's dissertational thesis [4]. Brown showed that the combinatorial structure of $\mathcal{SV}(\mathcal{U})$ is identical to the structure of the convex hull of $\mathcal{U}$. More precisely, a facet of the convex hull of $\mathcal{U}$ corresponds to an empty circle on the sphere, and therefore to a Voronoi vertex of $\mathcal{SV}(\mathcal{U})$. This immediately results in an $O(n \log n)$ time algorithm for the Voronoi diagram of points on the sphere.

Brown also considered the medial axis for a convex spherical region whose boundary is a simple closed curve consisting of geodesic arcs [4]. Another incremental algorithm for constructing the Voronoi diagram on the sphere was proposed by Augenbaum [2]. Recently, we developed an algorithm [14] to construct the Voronoi diagram of circular arcs (not necessarily geodesic) on the sphere, as a spherical analog of the method by Alt and Schwarzkopf [1].

Brown's most important discovery, however, is arguably the connection between 3-dimensional convex hulls and planar Voronoi diagrams [4, 5]. Let $\mathcal{U}$ be a set of points on the sphere, and consider a stereographic projection mapping the sphere to a plane. Let $\mathcal{W}$ be the resulting set of points in the plane. Brown showed that the Voronoi diagram of $\mathcal{W}$ (in the plane) has the same combinatorial structure as that part of the convex hull of $\mathcal{U}$ that is not visible from the projection center. This observation allowed Brown to compute the Voronoi diagram of a set of n points in the *plane* in time $O(n \log n)$.

The insight that Voronoi diagrams are so closely related to convex hulls considerably advanced our understanding of Voronoi diagrams at a time when computational geometry was in its infancy. Brown's transformation itself, however, never got much publicity. The lifting transformation, which uses a paraboloid instead of a sphere, quickly replaced it in the computational geometer's toolbox [17]. We will show in this paper that Brown's transformation is still a useful tool in computational geometry, and deserves not to be forgotten.

Brown used a stereographic projection to transform the (then) unfamiliar planar Voronoi diagram to a familiar convex hull in three dimensions. We employ it in the opposite direction, but with a similar goal: we transform the unfamiliar spherical Voronoi diagram of a set of sites on the sphere to (by now) familiar planar Voronoi diagrams of sites in the plane. We take advantage of the considerable work that has been done on the computation of planar Voronoi diagrams exactly and approximatively, not only in computational geometry, but also in CAGD [1, 6, 18]. We show that the spherical Voronoi diagram can be computed by computing two planar Voronoi diagrams and a little bit of glueing. We extend Brown's idea by using *two* stereographic projections (or inversions) with different projection centers on the sphere.

As a second result, we establish a surprising connection between planar furthest-site and nearest-site Voronoi diagrams. We show that a planar furthest-site Voronoi diagram can always be obtained as a portion of a nearest-site Voronoi diagram of a set of transformed sites.

As an immediate application of our transformations, we obtain an $O(n \log n)$ algorithm for computing the spherical Voronoi diagram of a set of circular arcs on the sphere, and for computing the planar furthest-site Voronoi diagram of a set of circular arcs in the plane.

2 Inversions and Voronoi diagrams

Let $\mathbb{E}^3$ be the 3-dimensional Euclidean vector space with the Euclidean norm $|\cdot|$. *Inversion* with the inversion center at the origin o is defined by $\zeta_o(v) = v/|v|^2$ [8]. More generally, the inversion

with inversion center at $\sigma \in \mathbb{E}^3$ is defined by $\zeta_\sigma(\nu) = \frac{\nu-\sigma}{|\nu-\sigma|^2} + \sigma$. This transformation has the following properties:

(1) The direction of the vector $\zeta_\sigma(\nu) - \sigma$ is the same as that of the vector $\nu - \sigma$, and the length of the vector $\zeta_\sigma(\nu) - \sigma$ is the inverse of the length of the vector $\nu - \sigma$.
(2) Inversion is involutory—application of inversion twice yields the original vector. We say that ν and $\zeta_\sigma(\nu)$ are inversive images of each other.
(3) The inversion ζ_σ maps a sphere not containing σ to a sphere not containing σ, a sphere containing σ to a plane not containing σ, and a plane containing σ to itself. If we consider planes to be spheres of infinite radius, we can say that ζ_σ maps spheres to spheres.

Fig. 1 shows two examples of inversion. The inversion in Fig. 1(b) maps the unit-diameter sphere $\mathcal{S}$ to the plane $\mathcal{T}$ and vice versa. Its restriction to $\mathcal{S}$ is known as the *stereographic projection from the origin* o. In the following, $\mathcal{S}$ will be a unit-diameter sphere, σ will be a point on $\mathcal{S}$, and $\mathcal{T}$ will be the plane tangent to $\mathcal{S}$ in the point opposite to σ. Therefore, $\mathcal{T}$ and $\mathcal{S}$ are inversive images of each other under ζ_σ. In Fig. 1(b), $\sigma = (0,0,0)$.

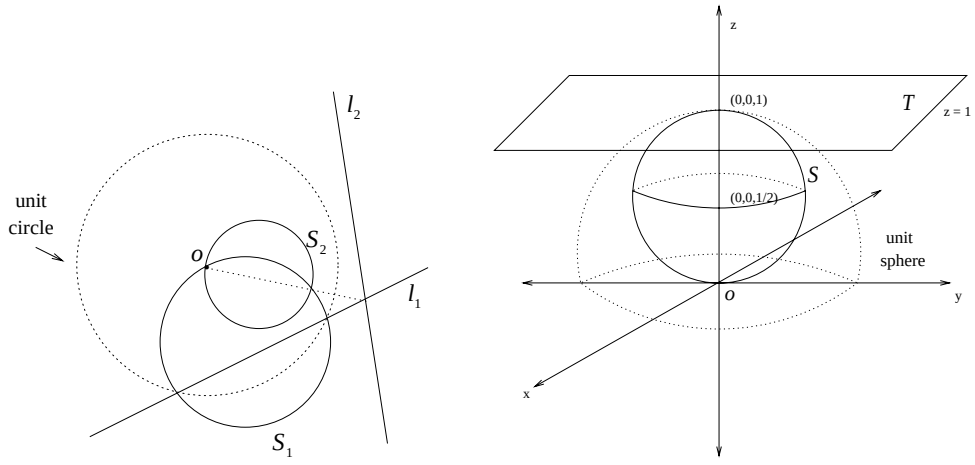

(a) two-dimensional example: line l_i and circle S_i, for $i = 1, 2$ are inversive pairs.

(b) three-dimensional example: sphere $\mathcal{S}$ and plane $\mathcal{T}$ is an inversive pair.

Figure 1: Examples of inversions.

Let a *site* be a compact subset of $\mathcal{T}$. We define the *near-type distance* d_n and the *far-type distance* d_f from a point x to a site w as follows:

$$d_n(x, w) := \inf_{y \in w} d(x, y),$$
$$d_f(x, w) := \sup_{y \in w} d(x, y).$$

Let w be a site in $\mathcal{T}$. Since w is compact, for each $x \in \mathcal{T}$ there exist $z_1, z_2 \in w$ such that $d_n(x, w) = d(x, z_1)$ and $d_f(x, w) = d(x, z_2)$. In particular, if w is a point, then $d_n(x, w) = d_f(x, w) = d(x, w)$. Furthermore, $d_n(\cdot, w), d_f(\cdot, w)$ are continuous functions defined in $\mathcal{T}$.

We denote the interior, the boundary, the closure, and the interior of the complement of a set A by $\operatorname{int} A$, ∂A, $\operatorname{cl} A$, and $\operatorname{ext} A$, respectively.

Definition 1 *Let $\mathcal{W}$ be a set of $n > 3$ disjoint sites in $\mathcal{T}$.*

For $w \in \mathcal{W}$, the nearest-site Voronoi region $VR(w,\mathcal{W})$ *and the* furthest-site Voronoi region $VR_f(w,\mathcal{W})$ *are defined as follows:*

$$\begin{aligned} VR(w,\mathcal{W}) &:= \{x \in \mathcal{T} \mid d_n(x,w) = \min_{w' \in \mathcal{W}} d_n(x,w')\}, \\ VR_f(w,\mathcal{W}) &:= \{x \in \mathcal{T} \mid d_f(x,w) = \max_{w' \in \mathcal{W}} d_f(x,w')\}. \end{aligned}$$

The (nearest-site) Voronoi diagram $\mathcal{V}(\mathcal{W})$ *is the tessellation $VR(w,\mathcal{W})_{w\in\mathcal{W}}$, the* furthest-site Voronoi diagram $\mathcal{V}_f(\mathcal{W})$ *is the tessellation $VR_f(w,\mathcal{W})_{w\in\mathcal{W}}$.*

The reader will be familiar with the near-type distance function d_n and the nearest-site Voronoi diagram it defines, but may wonder about the definition of the far-type distance d_f. The far-type distance is indeed the natural distance for defining furthest-site Voronoi diagrams, as we will see below.

The nearest-site Voronoi diagram defined above has been studied exhaustively in the literature, while the furthest-site diagram of sites (as opposed to points) does not seem to have received attention. We therefore now show that the furthest-site diagram is indeed a well-behaved tessellation of $\mathcal{T}$. In the following, $\mathcal{W}$ will always denote a set of $n > 3$ sites in $\mathcal{T}$.

Lemma 1 *$\mathcal{V}_f(\mathcal{W})$ is a tessellation of $\mathcal{T}$ satisfying*

(a) $\bigcup_{w\in\mathcal{W}} VR_f(w,\mathcal{W}) = \mathcal{T}$, *and*

(b) $VR_f(w,\mathcal{W}) \cap VR_f(w',\mathcal{W}) = \partial\, VR_f(w,\mathcal{W}) \cap \partial\, VR_f(w',\mathcal{W})$ *for $w, w' \in \mathcal{W}$.*

Proof. Let W be the union of all sites in $\mathcal{W}$.

(a) Let $x \in \mathcal{T}$. Since W is compact and $d(x,\cdot)$ is a continuous function, there exists a point $z \in W$ such that $d(x,z) = \max_{y\in W} d(x,y)$. Let w be the site containing z. Then $d(x,z) = d_f(x,w) = \max_{w'\in\mathcal{W}} d_f(x,w')$. Hence $x \in VR_f(w,\mathcal{W})$.

(b) For each $w, w' \in \mathcal{W}$ ($w \neq w'$), define $H(w,w') = \{y \in \mathcal{T} \mid d_f(y,w) \geqslant d_f(y,w')\}$. Since $d_f(\cdot,w)$ and $d_f(\cdot,w')$ are continuous, $H(w,w')$ is a closed subset of $\mathcal{T}$. Since $VR_f(w,\mathcal{W}) = \bigcap_{w'\in\mathcal{W}\setminus\{w\}} H(w,w')$, $VR_f(w,\mathcal{W})$ is a closed set. This implies that $\partial\, VR_f(w,\mathcal{W}) \cap \partial\, VR_f(w',\mathcal{W}) \subseteq VR_f(w,\mathcal{W}) \cap VR_f(w',\mathcal{W})$. If $x \in VR_f(w,\mathcal{W}) \cap VR_f(w',\mathcal{W})$ for distinct w, w', then there exists a disc D centered at x such that D contains $\mathcal{W}$ and its boundary passes through two points $z \in w, z' \in w'$. Consider a line l determined by x and z. Choose any point $x' \in l$, sufficiently close to x and lying in the opposite side of z as seen from x. Then z is the unique furthest point of W from x'. Hence $x' \notin VR_f(w',\mathcal{W})$ and thus $x \notin \operatorname{int} VR_f(w',\mathcal{W})$. Analogously we can prove that $x \notin \operatorname{int} VR_f(w,\mathcal{W})$. This shows that $VR_f(w,\mathcal{W}) \cap VR_f(w',\mathcal{W}) \subseteq \partial\, VR_f(w,\mathcal{W}) \cap \partial\, VR_f(w',\mathcal{W})$. □

For both diagrams, a *Voronoi vertex* v is a point which lies on the boundary of the Voronoi regions of at least three sites of $\mathcal{W}$, and a *Voronoi edge* e is a maximal connected subset of points lying on the boundary of the Voronoi regions for exactly two sites of $\mathcal{W}$.

In complete analogy to the above discussion, we can define sets $\mathcal{U}$ of n sites on the sphere $\mathcal{S}$, the near-type and far-type distance functions, Voronoi regions, and Voronoi diagrams. For clarity, we will call these the *spherical (nearest-site) Voronoi diagram* and and the *spherical furthest-site Voronoi diagram*, and denote them by $\mathcal{SV}(\mathcal{U})$ and $\mathcal{SV}_f(\mathcal{U})$.

Note that we can use either the geodesic distance on $\mathcal{S}$ or the Euclidean distance in 3 dimensions as the underlying distance of the near-type and far-type distance function. Both result in identical Voronoi diagrams.

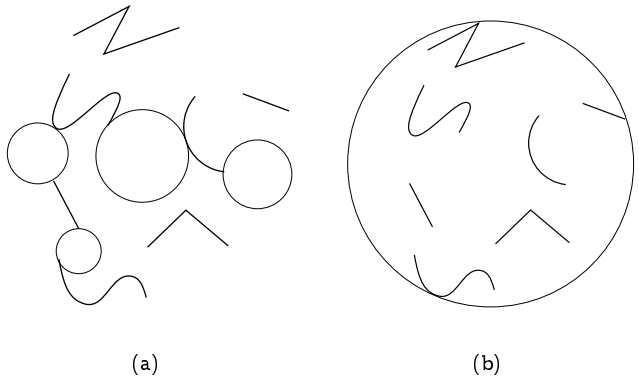

Figure 2: Largest empty discs (a) and smallest enclosing disc (b)

For a point $x \in \mathcal{T}$, we let $D(x)$ denote the *largest empty disc* with center x. It is the set of points $y \in \mathcal{T}$ such that $d(x,y) \leqslant \min_{w\in\mathcal{W}} d_n(x,w)$. Note that $D(x)$ contains no point of a site of $\mathcal{W}$ in its interior, but at least one point of a site on its boundary. Analogously, we let $D_f(x)$ denote the *smallest enclosing disc* with center x. It is the set of points $y \in \mathcal{T}$ such that $d(x,y) \leqslant \max_{w\in\mathcal{W}} d_f(x,w)$. Note that $D_f(x)$ contains all sites of $\mathcal{W}$, and at least one point of a site lies on its boundary. See Fig. 2.

Note that a point x is a vertex of $\mathcal{V}(\mathcal{W})$ if and only if $D(x)$ contains points of three or more different sites, and lies on an edge of $\mathcal{V}(\mathcal{W})$ if and only if $D(x)$ contains points of exactly two distinct sites. Analogously, x is a vertex of $\mathcal{V}_f(\mathcal{W})$ if and only if $D_f(x)$ has points of three or more different sites on its boundary, and lies on an edge of $\mathcal{V}_f(\mathcal{W})$ if and only if $D_f(x)$ has points of exactly two distinct sites on its boundary. This fundamental property relies on our definition of the far-type distance function — it would not be true if we had defined the furthest-site diagram using the near-type distance function.

We define the *spherical largest empty disc* $SD(x)$ and the *spherical smallest enclosing disc* $SD_f(x)$ analogously.

The following lemma shows that the spherical furthest-site Voronoi diagram is in fact also a spherical (nearest-site) Voronoi diagram. We denote the antipode (the diametrically opposite point) of a point $y \in \mathcal{S}$ by y^*, and the antipode of a site $u \subset \mathcal{S}$ by $u^* := \{y^* \mid y \in u\}$.

Lemma 2 (Brown [4]) *Let $\mathcal{U}$ be a set of sites on $\mathcal{S}$. Then $\mathcal{SV}_f(\mathcal{U}) = \mathcal{SV}(\mathcal{U}^*)$, where $\mathcal{U}^* = \{u^* \mid u \in \mathcal{U}\}$.*

Proof. For points $x, y \in \mathcal{S}$ we have $d(x,y) = \pi/2 - d(x,y^*)$. It follows that for a site $u \subset \mathcal{S}$, we have $d_f(x,u) = \pi/2 - d_n(x,u^*)$. Therefore, $x \in VR_f(u,\mathcal{U})$ if and only if $x \in VR(u^*,\mathcal{U}^*)$. The lemma follows. □

3 Convex hulls and Voronoi diagrams

Brown [4] observed that for a set $\mathcal{U}$ of n points on $\mathcal{S}$, the spherical Voronoi diagram $\mathcal{SV}(\mathcal{U})$ has the same combinatorial structure as the convex hull of $\mathcal{U}$. In fact, recall that a point $x \in \mathcal{S}$ is a vertex of $\mathcal{SV}(\mathcal{U})$ if and only if $SD(x)$ has three or more points on its boundary. Since $SD(x)$ is the intersection of $\mathcal{S}$ with a half-space, this implies the existence of an empy half-space containing

three points of $\mathcal{U}$ on its boundary. We thus have a one-to-one correspondence between Voronoi vertices and convex hull facets, and the same correspondence holds between Voronoi edges and convex hull edges.

Since the convex hull of n points in three dimensions can be computed in $O(n \log n)$ time, this allowed him to compute the spherical Voronoi diagram (and therefore also the spherical furthest-site Voronoi diagram) of a set of points within the same time bound.

Brown observed a second correspondence. Let $\mathcal{U}$ be a set of n points on $\mathcal{S} \setminus \{\sigma\}$, and let $\mathcal{W} := \zeta_\sigma(\mathcal{U})$ be its inversive image on $\mathcal{T}$. Consider a point $x \in \mathcal{S}$. Its largest empty disc $SD(x)$ contains a subset $U \subset \mathcal{U}$ of one or more points of $\mathcal{U}$ on its boundary.

If $SD(x)$ does not contain σ, then the inversive image $C := \zeta_\sigma(SD(x))$ is a disc in $\mathcal{T}$ not containing any point of $\mathcal{W}$. The points in $\zeta_\sigma(U)$ do lie on the boundary of C, and so C is the largest empty disc $D(y)$ of some point $y \in \mathcal{T}$ (note that in general $y \neq \zeta_\sigma(x)$). It follows that x is a vertex of $\mathcal{SV}(\mathcal{U})$ if and only if y is a vertex of $\mathcal{V}(\mathcal{W})$, and x lies on an edge of $\mathcal{SV}(\mathcal{U})$ if and only if y lies on an edge of $\mathcal{V}(\mathcal{W})$.

If $SD(x)$ contains σ in its interior, then the inversive image $\zeta_\sigma(SD(x))$ is the complement of an open disc C in $\mathcal{T}$, where C contains all the points of $\mathcal{W}$. Again the points in $\zeta_\sigma(U)$ lie on the boundary of C, and so C is the smallest enclosing disc $D_f(y)$ of some point $y \in \mathcal{T}$. It follows that x is a vertex of $\mathcal{SV}(\mathcal{U})$ if and only if y is a vertex of $\mathcal{V}_f(\mathcal{W})$, and x lies on an edge of $\mathcal{SV}(\mathcal{U})$ if and only if y lies on an edge of $\mathcal{V}_f(\mathcal{W})$.

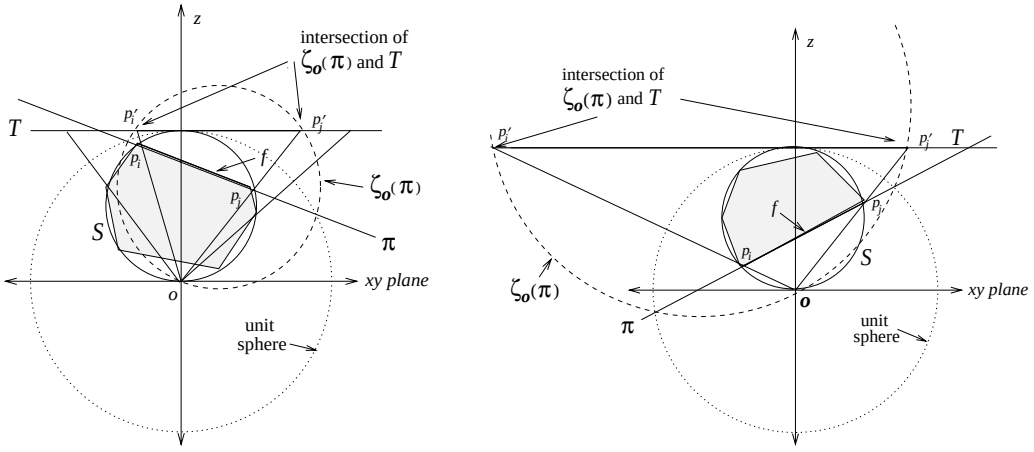

(a) f determines a vertex of $\mathcal{V}(\mathcal{W})$

(b) f determines a vertex of $\mathcal{V}_f(\mathcal{W})$

Figure 3: Illustration of the relation between a face f of convex hull of $\mathcal{U}$ and the planar Voronoi diagrams of $\mathcal{W} = \zeta_\sigma(\mathcal{U})$.

This observation, which precedes the well-known lifting transformation mapping the plane to a paraboloid, allowed Brown to compute the Voronoi diagram of a set of n points in the plane in time $O(n \log n)$. He transformed the points to the sphere by stereographic projection, computed their convex hull, and identified the facets pointing away from the origin.

We will go in the opposite direction. We wish to compute the spherical Voronoi diagram of a set of sites on the sphere. Reviving and generalizing Brown's results, we transform the spherical problem into a planar one, obtaining the spherical Voronoi diagram by piecing together two planar ones.

Note that while there is certainly again a relationship with the 3-dimensional convex hull of the sites, this does not appear to help us solve the problem, as little work has been done on the computation of convex hulls of curved objects in space.

4 Spherical and planar Voronoi diagrams

In the following, we assume again a set $\mathcal{U}$ of n sites on $\mathcal{S}$, and let $\mathcal{W} := \zeta_\sigma(\mathcal{U})$ be its inversive image on the plane $\mathcal{T}$. Generalizing Brown's result, we will show that part of $\mathcal{SV}(\mathcal{U})$ is isomorphic to $\mathcal{V}(\mathcal{W})$, while the other part is isomorphic to $\mathcal{V}_f(\mathcal{W})$.

The partition of $\mathcal{SV}(\mathcal{U})$ into two parts can be obtained by cutting the sphere $\mathcal{S}$ with a closed curve, which we obtain as follows. Let

$$\mathcal{Z}_\mathcal{S}(\sigma,\mathcal{U}) = \{x \in \mathcal{S} \mid d(x,\sigma) \leqslant d_n(x,u) \text{ for all } u \in \mathcal{U}\}.$$

So $\mathcal{Z}_\mathcal{S}(\sigma,\mathcal{U})$ is in fact the Voronoi region of σ, if σ was added to the set $\mathcal{U}$ as a point site.

We now partition the sphere $\mathcal{S}$ into three disjoint regions:

$$\begin{aligned}
\mathcal{R}_\sigma^+ &:= \operatorname{ext} \mathcal{Z}_\mathcal{S}(\sigma,\mathcal{U}) = \{x \in \mathcal{S} \mid \sigma \in \operatorname{ext} SD(x)\},\\
\mathcal{R}_\sigma^- &:= \operatorname{int} \mathcal{Z}_\mathcal{S}(\sigma,\mathcal{U}) = \{x \in \mathcal{S} \mid \sigma \in \operatorname{int} SD(x)\},\\
\overline{\mathcal{R}}_\sigma &:= \partial\mathcal{Z}_\mathcal{S}(\sigma,\mathcal{U}) = \{x \in \mathcal{S} \mid \sigma \in \partial SD(x)\}.
\end{aligned}$$

The region $\mathcal{R}_\sigma^-$ is star-shaped from σ, and so $\overline{\mathcal{R}}_\sigma$ is a closed Jordan curve separating $\mathcal{R}_\sigma^+$ and $\mathcal{R}_\sigma^-$.

We now define two functions as follows:

$$\begin{aligned}
\phi_\sigma:&\quad \mathcal{R}_\sigma^+ \mapsto \mathcal{T} \quad \phi_\sigma(x) := y \in \mathcal{T} \text{ such that } D(y) = \zeta_\sigma(SD(x))\\
\psi_\sigma:&\quad \mathcal{R}_\sigma^- \mapsto \mathcal{T} \quad \psi_\sigma(x) := y \in \mathcal{T} \text{ such that } \operatorname{ext} D_f(y) = \operatorname{int} \zeta_\sigma(SD(x))
\end{aligned}$$

Figs. 4 and 5 illustrate the two definitions. Note that both functions depend on the set $\mathcal{U}$.

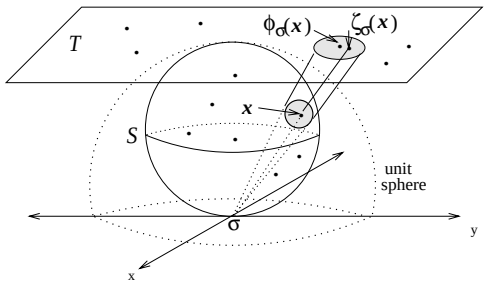

Figure 4: The definition of ϕ_σ.

Lemma 3 *The functions ϕ_σ and ψ_σ are well-defined, bijective, and bicontinuous.*

Proof. Let $x \in \mathcal{R}_\sigma^+$. Then the inversive image $\zeta_\sigma(SD(x))$ is a closed disc whose interior does not intersect $\mathcal{W}$, and so there is a unique point $y \in \mathcal{T}$ with $D(y) = \zeta_\sigma(SD(x))$. On the other hand, for any point $y \in \mathcal{T}$, the inversive image $\zeta_\sigma(D(y))$ is an empty spherical disc not containing σ and so there is $x \in \mathcal{R}_\sigma^+$ with $\zeta_\sigma(SD(x)) = D(y)$. It follows that ϕ_σ is a bijective function. Bicontinuity follows from the continuity of $d_n(\cdot, w)$.

Let now $x \in \mathcal{R}_\sigma^-$. Then the inversive image $\zeta_\sigma(SD(x))$ is the exterior of an open disc, and so there is $y \in \mathcal{T}$ with $\operatorname{ext} D_f(y) = \operatorname{int} \zeta_\sigma(SD(x))$. For any $y \in \mathcal{T}$, the inversive image $\zeta_\sigma(\operatorname{ext} D_f(y))$ is

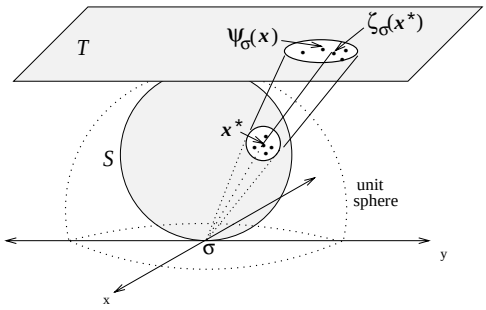

Figure 5: The definition of ψ_σ.

a spherical disc containing σ in its interior, and so there is $x \in \mathcal{R}_\sigma^-$ with $\operatorname{int} \zeta_\sigma(\mathrm{SD}(x)) = \operatorname{ext} \mathrm{D}_f(y)$. So ψ_σ is bijective, and bicontinuity follows from the continuity of $d_n(\cdot, u)$ and $d_f(\cdot, w)$. □

Lemma 4 *Let $x \in \mathcal{R}_\sigma^+$ and $u \in \mathcal{U}$. Then u intersects $\partial\mathrm{SD}(x)$ if and only if $\zeta_\sigma(u)$ intersects $\partial\mathrm{D}(\phi_\sigma(x))$.*

Let $x \in \mathcal{R}_\sigma^-$ and $u \in \mathcal{U}$. Then u intersects $\partial\mathrm{SD}(x)$ if and only if $\zeta_\sigma(u)$ intersects $\partial\mathrm{D}_f(\psi_\sigma(x))$.

Proof. Let $x \in \mathcal{R}_\sigma^+$. We have

$$\partial\mathrm{D}(\phi_\sigma(x)) \cap \zeta_\sigma(u) = \zeta_\sigma(\partial\mathrm{SD}(x)) \cap \zeta_\sigma(u) = \zeta_\sigma(\partial\mathrm{SD}(x) \cap u).$$

For $x \in \mathcal{R}_\sigma^-$ we have

$$\partial\mathrm{D}_f(\psi_\sigma(x)) \cap \zeta_\sigma(u) = \zeta_\sigma(\partial\mathrm{SD}(x)) \cap \zeta_\sigma(u) = \zeta_\sigma(\partial\mathrm{SD}(x) \cap u).$$

The lemma follows. □

Corollary 1 *A point $x \in \mathcal{R}_\sigma^+$ is a vertex, lies on an edge, or lies in a cell of $\mathcal{SV}(\mathcal{U})$ if and only if $\phi_\sigma(x)$ is a vertex, lies on an edge, or lies in a cell of $\mathcal{V}(\mathcal{W})$.*

A point $x \in \mathcal{R}_\sigma^-$ is a vertex, lies on an edge, or lies in a cell of $\mathcal{SV}(\mathcal{U})$ if and only if $\psi_\sigma(x)$ is a vertex, lies on an edge, or lies in a cell of $\mathcal{V}_f(\mathcal{W})$.

In other words, the part of $\mathcal{SV}(\mathcal{U})$ lying in $\mathcal{R}_\sigma^+$ is the continuous image of $\mathcal{V}(\mathcal{W})$, while the part of $\mathcal{SV}(\mathcal{U})$ lying in $\mathcal{R}_\sigma^-$ is the continuous image of $\mathcal{V}_f(\mathcal{W})$.

To summarize, the spherical Voronoi diagram $\mathcal{SV}(\mathcal{U})$ is partitioned by the curve $\overline{\mathcal{R}}_\sigma$. The two regions $\mathcal{R}_\sigma^+$ and $\mathcal{R}_\sigma^-$ can be identified with the plane $\mathcal{T}$, and the parts of $\mathcal{SV}(\mathcal{U})$ in the two regions are homeomorphic to $\mathcal{V}(\mathcal{W})$ and $\mathcal{V}_f(\mathcal{W})$. The curve $\overline{\mathcal{R}}_\sigma$ itself corresponds to the "points at infinity" in both planes. We study the intersection of $\mathcal{SV}(\mathcal{U})$ with $\overline{\mathcal{R}}_\sigma$ in a bit more detail.

Lemma 5 *An edge of $\mathcal{SV}(\mathcal{U})$ intersects $\overline{\mathcal{R}}_\sigma$ in at most two points. It follows that the intersection of $\mathcal{SV}(\mathcal{U})$ with $\overline{\mathcal{R}}_\sigma$ is a set of $O(n)$ isolated points, where $n = |\mathcal{U}|$.*

Proof. Assume there were three points x_1, x_2, x_3, all in $\overline{\mathcal{R}}_\sigma$, and all on the same edge of $\mathcal{SV}(\mathcal{U})$ separating the Voronoi regions of $u, u' \in \mathcal{U}$. Consider the empty discs $\mathrm{SD}(x_i)$, $i = 1, 2, 3$. We have $o \in \bigcap_i \mathrm{SD}(x_i)$, and $\left(\operatorname{int} \bigcup_i \mathrm{SD}(x_i)\right) \cap (u \cup u') = \emptyset$. On the other hand, there are points $u_i \in u \cap \partial\mathrm{SD}(x_i)$ and $u_i' \in u' \cap \partial\mathrm{SD}(x_i)$, for $i = 1, 2, 3$. We can now construct a planar embedding of the complete bipartite graph $K_{3,3}$: The nodes on one side are x_1, x_2, x_3, the nodes on the other side are o, u, u'. Graph edges are the segments $x_i o$, $x_i u_i$, and $x_i u_i'$. Since $K_{3,3}$ is not planar, we

obtain a contradiction proving the lemma. □

This leads to the following procedure to compute $\mathcal{SV}(\mathcal{U})$: Compute the inversive image $\mathcal{W}$ of $\mathcal{U}$, and the planar Voronoi diagrams $\mathcal{V}(\mathcal{W})$ and $\mathcal{V}_f(\mathcal{W})$. Map these diagrams back to the sphere using ϕ_σ^{-1} and ψ_σ^{-1}, and identify the $O(n)$ endpoints of edges in both diagrams lying on $\overline{\mathcal{R}}_\sigma$ (these are the "endpoints at infinity" of the semi-infinite edges of the two planar diagrams). Sort these endpoints along $\overline{\mathcal{R}}_\sigma$, and merge identical ones from the two parts of the diagram, resulting in $\mathcal{SV}(\mathcal{U})$.

The only problem with this approach is that it assumes that we know how to compute $\mathcal{V}_f(\mathcal{W})$. While nearest-site Voronoi diagrams in the plane have been studied in much detail, this is not true for furthest-site diagrams. In the next section we therefore give a method to compute $\mathcal{SV}(\mathcal{U})$ using only *nearest-site* planar diagrams.

5 Two inversions suffice

We will compute the spherical Voronoi diagram of a set of n sites $\mathcal{U}$ on $\mathcal{S}$, assuming that we have a suitable algorithm $\mathcal{A}$ for computing planar Voronoi diagrams. We will not consider the meaning of "suitable" here in more detail—it depends on the nature of the sites $\mathcal{U}$. If, for instance, $\mathcal{U}$ is a set of circular arcs on $\mathcal{S}$, then $\mathcal{A}$ needs to be able to compute the planar Voronoi diagram of a set of circular arcs in the plane.

We have seen in the previous section that the portion of $\mathcal{SV}(\mathcal{U})$ lying in $\mathcal{R}_\sigma^-$ corresponds directly to the planar Voronoi diagram $\mathcal{V}(\zeta_\sigma(\mathcal{U}))$. To compute $\mathcal{SV}(\mathcal{U})$, we still need to fill in the part of $\mathcal{SV}(\mathcal{U})$ lying in $\mathcal{R}_\sigma^-$. Our key idea is to apply a *second inversion*, using a different inversion center.

Lemma 6 *Let $\mathcal{U}$ be a set of $n > 3$ sites on $\mathcal{S}$, and let $\mathcal{W} := \zeta_\sigma(\mathcal{U})$ be its inversive image on $\mathcal{T}$.*

There is a point $\eta \in \mathcal{S}$ not in any site of $\mathcal{U}$, such that $\zeta_\sigma(\eta)$ lies in the interior of the convex hull of $\mathcal{W}$. We have

$$\mathrm{cl}\,\mathcal{R}_\sigma^- \subset \mathcal{R}_\eta^+.$$

Proof. Since $n > 3$, there must be a point $q \in \mathcal{T}$ in the interior of the convex hull of $\mathcal{W}$ not contained in any site of $\mathcal{W}$. Let $\eta := \zeta_\sigma(q)$.

Consider now a point $x \in \mathrm{cl}\,\mathcal{R}_\sigma^- = \mathcal{R}_\sigma^- \cup \overline{\mathcal{R}}_\sigma$. Then $\sigma \in \mathrm{SD}(x)$, and so the inversive image $\zeta_\sigma(\mathrm{SD}(x))$ is the complement of an open circle or open half-plane containing $\mathcal{W}$. Since $q = \zeta_\sigma(\eta)$ is contained in the interior of the convex hull of $\mathcal{W}$, $\zeta_\sigma(\eta) \notin \zeta_\sigma(\mathrm{SD}(x))$, and so $\eta \notin \mathrm{SD}(x)$. Therefore $x \in \mathcal{R}_\eta^+$. □

We will use a second inversion ζ_η with inversion center η, where η is as in Lemma 6. It maps the sphere $\mathcal{S}$ to the plane $\mathcal{T}'$ tangent to $\mathcal{S}$ in the antipode of η. Let $\mathcal{W}'$ be the inversive image of $\mathcal{U}$ under ζ_η, and let $\sigma' := \zeta_\eta(\sigma)$.

Lemma 7 *The function ϕ_η is defined for all points in $\mathrm{cl}\,\mathcal{R}_\sigma^-$, and we have*

$$\phi_\eta(\mathrm{cl}\,\mathcal{R}_\sigma^-) = \mathcal{Z}_{\mathcal{T}'}(\sigma', \mathcal{W}').$$

Proof. By Lemma 6 we have $\mathrm{cl}\,\mathcal{R}_\sigma^- \subset \mathcal{R}_\eta^+$ and so ϕ_σ is defined for all points in $\mathrm{cl}\,\mathcal{R}_\sigma^-$. Let now $x \in \mathcal{R}_\eta^+$. We have $\zeta_\eta(\mathrm{SD}(x)) = \mathrm{D}(\phi_\eta(x))$ (this is an empty disc on $\mathcal{T}'$), and so the following

statements are all equivalent:

$$\begin{aligned} \sigma &\in \mathrm{SD}(x) \\ \sigma' = \zeta_\eta(\sigma) &\in \zeta_\eta(\mathrm{SD}(x)) = \mathrm{D}(\phi_\eta(x)) \\ \phi_\eta(x) &\in \mathcal{Z}_{\mathcal{T}'}(\sigma', \mathcal{W}'), \end{aligned}$$

which proves the lemma. □

Lemma 8 *The portion of $\mathcal{SV}(\mathcal{U})$ contained in $\mathcal{R}_\sigma^-$ is homeomorphic to the portion of $\mathcal{V}(\mathcal{W}')$ contained in $\mathcal{Z}_{\mathcal{T}'}(\sigma', \mathcal{W}')$.*

Proof. By Lemma 6, $\mathcal{R}_\sigma^-$ is contained in $\mathcal{R}_\eta^+$. The portion of $\mathcal{SV}(\mathcal{U})$ contained in $\mathcal{R}_\eta^+$ is homeomorphic under ϕ_η to $\mathcal{V}(\mathcal{W}')$ in $\mathcal{T}'$ by Cor. 1. By Lemma 7, ϕ_η maps $\mathcal{R}_\sigma^-$ to $\mathcal{Z}_{\mathcal{T}'}(\sigma', \mathcal{W}')$, so the portion of $\mathcal{SV}(\mathcal{U})$ contained in $\mathcal{R}_\sigma^-$ is homeomorphic to the portion of $\mathcal{V}(\mathcal{W}')$ contained in $\mathcal{Z}_{\mathcal{T}'}(\sigma', \mathcal{W}')$. □

We now give our algorithm to compute the spherical Voronoi diagram $\mathcal{SV}(\mathcal{U})$ of a set of $n > 3$ sites $\mathcal{U}$. We first choose a point $\sigma \in \mathcal{S}$ such that it does not lie in any site. We compute the inversive image $\mathcal{W} := \zeta_\sigma(\mathcal{U})$ of $\mathcal{U}$, apply a planar Voronoi diagram algorithm, and obtain $\mathcal{V}(\mathcal{W})$. We now pick a point q in the interior of the convex hull of $\mathcal{W}$ and not in any site of $\mathcal{W}$. Let $\eta := \zeta_\sigma(q)$ be its inversive image, and let ζ_η be the inversion with inversion center η. Let $\mathcal{W}' := \zeta_\eta(\mathcal{U})$ be the inversive image of $\mathcal{U}$. We apply again a planar Voronoi diagram algorithm to compute $\mathcal{V}(\mathcal{W}')$ (a Voronoi diagram on $\mathcal{T}'$). We identify the portion of $\mathcal{V}(\mathcal{W}')$ lying in $\mathcal{Z}_{\mathcal{T}'}(\zeta_\eta(\sigma), \mathcal{W}')$. This can be done by traversing the diagram and testing the distance of each vertex or edge to $\zeta_\eta(\sigma)$.

Finally, we map $\mathcal{V}(\mathcal{W})$ and the portion of $\mathcal{V}(\mathcal{W}')$ to $\mathcal{S}$ through ϕ_σ^{-1} and ϕ_η^{-1}. We identify the $O(n)$ endpoints on $\overline{\mathcal{R}}_\sigma$. These are the "endpoints at infinity" of $\mathcal{V}(\mathcal{W})$, and the points of $\mathcal{V}(\mathcal{W}')$ lying on the boundary of $\mathcal{Z}_{\mathcal{T}'}(\zeta_\eta(\sigma), \mathcal{W}')$. We sort these endpoints along $\overline{\mathcal{R}}_\sigma$, and merge identical ones, resulting in $\mathcal{SV}(\mathcal{U})$.

Theorem 1 *Given an algorithm $\mathcal{A}$ to compute the Euclidean Voronoi diagram of a set of sites in the plane, we can compute the Voronoi diagram of a set of n sites on the sphere by two applications of $\mathcal{A}$ to sets of n transformed sites, and $O(n \log n)$ extra computation involving primitive operations on sites.*

Using Fortune's [10] or Alt and Schwarzkopf's [1] algorithm for the Voronoi diagram of circular arcs in the plane, we have the following corollary.

Corollary 2 *The spherical Voronoi diagram of n circular arcs (not necessarily geodesic) on the sphere can be computed in $O(n \log n)$ time.*

6 Computing the furthest-site Voronoi diagram

By Lemma 2 the spherical furthest-site Voronoi diagram of a set of sites is nothing else but the nearest-site Voronoi diagram of the sites' antipodes. We now show, a little surprising, that a *planar* furthest-site Voronoi diagram can also be obtained using an algorithm for nearest-site diagrams.

Given a set $\mathcal{W}$ of n sites in the plane $\mathcal{T}$, let $\mathcal{S}$ be tangent to $\mathcal{T}$, and σ the inversion center that maps $\mathcal{T}$ to $\mathcal{S}$ and vice versa. We pick a point q in the interior of the convex hull of $\mathcal{W}$ but not in any site. Let $\mathcal{U} := \zeta_\sigma(\mathcal{W})$ be the inversive image of $\mathcal{W}$, let $\eta := \zeta_\sigma(q)$ be the inversive

image of q, let ζ_η be the inversion with inversion center η, and let $\mathcal{W}' := \zeta_\eta(\mathcal{U}) = \zeta_\eta(\zeta_\sigma(\mathcal{W}))$. By Lemma 4, $\mathcal{V}_f(\mathcal{W})$ is homeomorphic under ψ_σ to the portion of $\mathcal{SV}(\mathcal{U})$ contained in $\mathcal{R}_\sigma^-$. On the other hand, this portion of $\mathcal{SV}(\mathcal{U})$ is homeomorphic under ϕ_η to the portion of $\mathcal{V}(\mathcal{W}')$ contained in $\mathcal{Z}_{\mathcal{T}'}(\zeta_\eta(\sigma), \mathcal{W}')$.

We can therefore compute $\mathcal{V}_f(\mathcal{W})$ by computing $\mathcal{W}'$ through two inversions, computing $\mathcal{V}(\mathcal{W}')$, identifying the portion contained in $\mathcal{Z}_{\mathcal{T}'}(\zeta_\eta(\sigma), \mathcal{W}')$, and mapping this portion back to $\mathcal{T}$ through ϕ_η^{-1} and ψ_σ.

Theorem 2 *Given an algorithm $\mathcal{A}$ to compute the Euclidean Voronoi diagram of a set of sites in the plane, we can compute the furthest-point Voronoi diagram of a set of n sites in the plane by one application of $\mathcal{A}$ on a set of n transformed sites, and $O(n)$ extra computation involving primitive operations on sites.*

Again we formulate a specific result for circular arcs.

Corollary 3 *The furthest-site Voronoi diagram of n circular arcs in the plane can be computed in $O(n \log n)$ time.*

References

[1] H. Alt and O. Schwarzkopf. The Voronoi diagram of curved objects. In *Proc. 11th Annu. ACM Symposium on Computation Geometry* (1995), pp. 89-97.

[2] J. M. Augenbaum and Charles S. Peskin. On the construction of the Voronoi mesh on the sphere. *Journal of Computational Physics* 59 (1985), pp. 177-192.

[3] F. Aurenhammer. Voronoi diagrams : a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23 (1991), pp. 345-405.

[4] K. Q. Brown. *Geometric transforms for fast geometric algorithms.* Ph.D. dissertation, Carnegie-Mellon Univ., 1980, pp. 73-75.

[5] K. Q. Brown. Voronoi diagrams from convex hulls. *Information Processing Letters* 9 (1979), pp. 223–228.

[6] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific J. Math.* 181 (1997), pp. 57-88.

[7] H. I. Choi, S. W. Choi, H. P. Moon, and N. S. Wee. New algorithm for medial axis transform of plane domain, *Graphical Models and Image Processing* 59 (1997), pp. 463-483.

[8] C. W. Dodge. *Euclidean Geometry and Transformations.* Addison-Wesley, 1972.

[9] R. T. Farouki and J. K. Johnstone. The bisector of a point and a plane parametric curve. *Comput. Aided Geom. Design* 11 (1994), pp. 117-151.

[10] S. J. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica* 2 (1987), pp. 153–174.

[11] R. Klein. *Concrete and Abstract Voronoi Diagrams.* Vol. 400 of LNCS, Springer-Verlag, 1989.

[12] R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract Voronoi Diagrams. *Computational Geometry : Theory and Applications* 3 (1993), pp. 157-184.

[13] J. R. Munkres. *Topology*. Prentice-Hall, 1975, pp. 385.

[14] H.S. Na, C. N. Lee, and O. Cheong. Randomized incremental construction of spherical Voronoi diagram for points and circular arcs. *Manuscript*, 1998.

[15] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, 1992.

[16] B. O'Neill. *Elementary Differential geometry*. 2nd edition, Academic Press, 1966, pp. 319-320.

[17] F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985, pp. 15-17.

[18] R. Ramamurthy and R. T. Farouki. Voronoi diagram and medial axis algorithm for planar domains with curved boundaries: I. Theoretical foundations. *Journal of Computational and Applied Mathematics* 102 (1999), pp. 117.

[19] R. Ramamurthy and R. T. Farouki. Specified-precision computation of curve/curve bisectors. *Internat. J. Comput. Geom. Appl.* 8 (1998), pp. 599–619.

[20] T. Rausch, F. E. Wolter, and O. Sniehotta. Computation of medial curves on surfaces. *The Mathematics of Surfaces* VII (1997), pp. 43-68.