Title
Realizing Farthest-Point Voronoi Diagrams.

Authors
Biedl, Therese C
Grimm, Carsten
Palios, Leonidas
et al.

Publication Date
2016

Peer reviewed

Realizing Farthest-Point Voronoi Diagrams

Therese Biedl * Carsten Grimm †‡ Leonidas Palios § Jonathan Shewchuk ¶ Sander Verdonschot ‖

Abstract

The farthest-point Voronoi diagram of a set of n sites is a tree with n leaves. We investigate whether arbitrary trees can be realized as farthest-point Voronoi diagrams. Given an abstract ordered tree T with n leaves and prescribed edge lengths, we produce a set of n sites S in $O(n)$ time such that the farthest-point Voronoi diagram of S represents T. We generalize this algorithm to smooth, strictly convex, symmetric distance functions. Lastly, given a subdivision Z of $\mathbb{R}^k$ with k a small constant, we check in linear time whether Z realizes a k-dimensional farthest-point Voronoi diagram.

1 Background

In 1999, Liotta and Meijer posed the following question: Given a tree T, can one draw T in the plane so that the resulting embedding is the Voronoi diagram of some set of sites in the plane? They consider the *ordered model*: The tree T is given as an abstract ordered tree, i.e., as a set of vertices, a set of edges, and a cyclic order of the of the edges incident to each vertex. We are searching for a set of sites S such that the vertices and edges of the Voronoi diagram of S form an embedding of T that respects the cyclic order of the edges around each vertex in T. Liotta and Meijer showed that every ordered tree can be realized as a Voronoi diagram [7, 8].

Quite related to this is the *Inverse Voronoi Problem*, which asks the question in the *geometric model*. Here we are given a tree (or more generally a graph) and also a drawing of it, i.e., coordinates for all interior nodes and rays to infinity for all edges to leaves. We are searching for a set of sites S such that the Voronoi diagram of S is exactly this tree with this drawing. The problem was introduced by Ash and Bolker [4] and the question can be answered in linear time [6], even if the tree has vertices of degree exceeding three [5].

A number of variants have been studied. Aloupis et al. [3] posed an extension-version of the Inverse Voronoi Problem. Other papers study the straight skeleton, rather than the Voronoi diagram. Aichholzer et al. resolved this for the ordered model [2], and (with different coauthors) for the ordered model where edge directions are given [1]. The Inverse Straight Skeleton Problem was resolved by Biedl et al. [5].

Our results. We ask whether trees can be realized by yet another computational geometry construct, namely, the *farthest-point Voronoi diagram* (defined below). We consider both models and obtain the following results.

Ordered Model: Similarly as in [3, 8], for the ordered model the answer is always "yes". Thus for any given ordered tree T, we can find a set of sites S in convex position such that the farthest-point Voronoi diagram of S is T, with the edges in the specified order. In contrast to related results, we can also realize edge lengths, i.e., if each interior edge e is assigned a positive weight $w(e)$, then we can find sites so that e has length $w(e)$.

We give the construction first for the "normal" (Euclidean) farthest-point Voronoi diagram, and then generalize it to any convex distance function for which the unit circle is smooth and strictly convex.

Geometric Model: Similarly as in [5, 6], for the geometric model not every geometric tree can be realized. Nonetheless, one can test in polynomial time whether for a given geometric tree T there exists a set of points whose farthest-point Voronoi diagram is T. If so, then the set of sites is not always unique, but it can be described as the solution space of a linear program.

We describe this result for arbitrary fixed dimension. For a given convex subdivision Z of $\mathbb{R}^k$ with n cells, we formulate a linear program with k variables that tests whether there exists a set of n sites whose farthest-point Voronoi diagram realizes Z. This linear program can be solved in linear time if k is a small constant [10].

2 Preliminaries

Let S be a set of sites and let p be a point in the plane. Let $\mathrm{F}_S(p)$ be the smallest disc centered at p that con-

*David R. Cheriton School of Computer Science, University of Waterloo, Canada, biedl@uwaterloo.ca. Supported by NSERC.

†Computational Geometry Lab, School of Computer Science, Carleton University, Ottawa, Ontario, Canada.

‡Institut für Simulation und Graphik, Fakultät für Informatik, Otto-von-Guericke-Universität Magdeburg, Magdeburg, Germany carsten.grimm@ovgu.de.

§Dept. of Computer Science and Engineering, University of Ioannina, Greece, palios@cse.uoi.gr.

¶Computer Science Division, University of California, Berkeley, CA, USA, jrs@cs.berkeley.edu. Supported by the National Science Foundation under Award CCF-1423560.

‖School of Electrical Engineering and Computer Science, University of Ottawa, Ottawa, ON, Canada, sander@cg.scs.carleton.ca. Supported by NSERC and the Ontario Ministry of Research and Innovation.

tains all sites in S; we call this the *full disc* of p with respect to S. For a set S of sites, the *farthest-point Voronoi diagram* of S, denoted by fVor(S), is defined as follows: A point p is a vertex of fVor(S) if and only if $\mathrm{F}_S(p)$ passes through three or more sites in S. A point p is located in the relative interior of an edge of fVor(S) if and only if $\mathrm{F}_S(p)$ passes through exactly two sites in S. fVor(S) divides the plane into convex cells, and one easily verifies that each cell consists of all points that are farthest from one site s. We say that site s is *relevant* if there is a point in the plane for which s is a farthest point, and *proper* if there is a point for which s is the *unique* farthest point. (For strictly convex distance functions "relevant" and "proper" are the same thing; see Section 4.2 for more details.)

The structure of the farthest-point Voronoi diagram is closely related to the convex hull CH(S) of S: (i) A site $s \in S$ is proper if and only if s is an extreme point of S. (ii) Two sites s and s' are adjacent along CH(S) if and only if the farthest-point Voronoi cells of s and s' share an unbounded edge (ray or line). (iii) The circular order of the sites along CH(S) is the circular order of the farthest-point Voronoi cells in fVor(S).

3 Ordered Trees

Consider the farthest-point Voronoi diagram fVor(S) of a set S of sites in the plane. We introduce symbolic vertices as endpoints for the unbounded edges of fVor(S). We say that fVor(S) is a *realization* of an ordered tree T if T is isomorphic to the abstract ordered tree formed by the Voronoi vertices, the symbolic vertices and the Voronoi edges of fVor(S). In the following, we consider only ordered trees without degree two vertices, since there are no degree two vertices in a farthest-point Voronoi diagram.

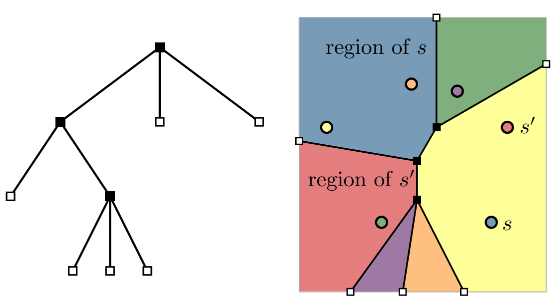

(a) (b)

Figure 1: (a) An ordered tree T; (b) a realization of T as a farthest-point Voronoi diagram. Empty squares are leaves; also symbolic endpoints of unbounded edges.

Given an ordered tree T, we seek to determine a set S of sites in the plane such that fVor(S) realizes T. We proceed in an incremental fashion where we place sites to create the internal vertices of T one by one.

Realizing a star. We begin with an ordered tree T_1 with one internal node v of degree ℓ. We realize T_1 by placing ℓ sites $s_1, s_2, \ldots, s_\ell$ on a unit circle C centered at the origin. The origin becomes the Voronoi vertex that we identify with v.

Any subsequent site s has to be placed at a location that is *safe* for the current sites S in the following sense: Every vertex in the diagram for S remains a vertex in the diagram for $S \cup \{s\}$ and every bounded edge in the diagram for S remains a bounded edge in the diagram for $S \cup \{s\}$. It is acceptable for a safe site to increase the degree of a vertex of the diagram. After the initial step, every point s strictly inside C is safe.[1]

On the other hand, any subsequent site s must be proper. Any site outside the convex hull CH(S) is proper.[1] Thus, all additional sites will be placed in the *lunes* that remain when we remove CH($\{s_1, \ldots, s_\ell\}$) from the disc bounded by C.

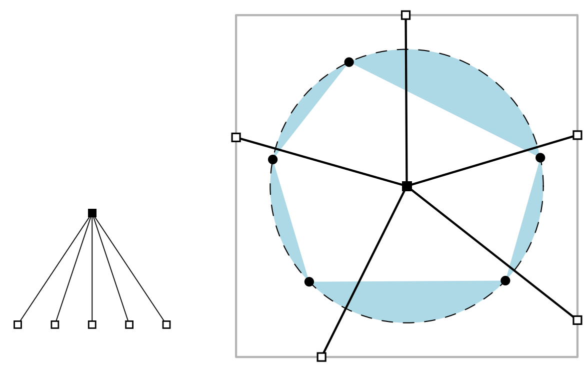

(a) (b)

Figure 2: (a) An ordered tree with one internal vertex; (b) a realization of that ordered tree as a farthest-point Voronoi diagram. All subsequent sites will be placed in the lunes (shaded blue).

Realizing larger trees. Suppose we can realize every ordered tree with $k \geq 1$ internal vertices as farthest-point Voronoi diagram, for some $k \in \mathbb{N}$. Consider an ordered tree T_{k+1} with $k+1$ internal vertices. There is an internal vertex v in T_{k+1} that becomes a leaf when all leaves adjacent to v are deleted. Let T_k be the tree that results from deleting the leaves adjacent to v. Since T_k is an ordered tree with k internal vertices, we can find a set S of sites such that fVor(S) is a realization of T_k. We seek to place additional sites such that the resulting farthest-point Voronoi diagram realizes T_{k+1}.

Vertex v is a leaf in T_k, hence corresponds to a symbolic endpoint in fVor(S) that lies on a ray r. Let u be the internal vertex at which r ends (hence u is the neighbor of v in T_k). Ray r separates the regions of two

[1]In the appendix, we provide full proofs for the claim for smooth, strictly convex, symmetric distance functions.

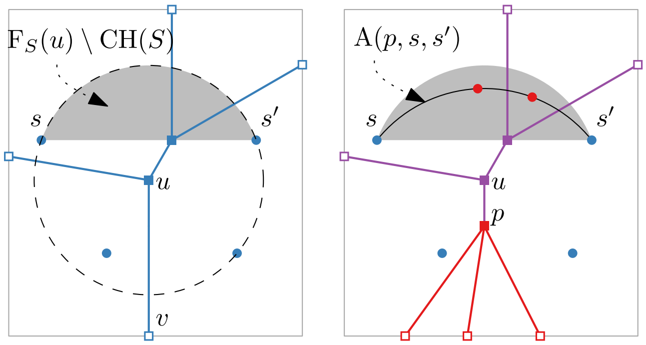

Figure 3: Extending the realization of an ordered tree.

sites s and s', so by the definition of $\mathrm{fVor}(S)$, for every point $p \in r$ the full disc $F_S(p)$ goes through s and s' and contains all other sites in its interior.

We want to place sites such that we create a Voronoi vertex at some point p on ray r (and then assign this point to v). To create a Voronoi vertex at p, we have to place a new site s'' on the boundary of $F_S(p)$. To make its region appear between the ones of s and s', we should place s'' on the (shorter) circular arc $A(p, s, s')$ from s to s' along $F_S(p)$. If v is adjacent to ℓ leaves in T_{k+1} ($\ell > 1$ since we have no vertices of degree 2), then we should place $\ell - 1$ new sites along $A(p, s, s')$.

Observe that the choice of p is arbitrary, as long as it is on the ray. We can therefore choose the distance between u and p (the future location of v) and realize any specified edge length of (u, v). To summarize, we can realize every ordered tree T as a farthest-point Voronoi diagram by placing the sites for some vertex of T on a circle and then repeatedly expanding the resulting farthest-point Voronoi diagram by placing the next vertex on the appropriate ray and sites for it on the corresponding arc. We place n sites for an ordered tree with n leaves. The entire construction takes $O(n)$ time, since computing the coordinates of each site takes constant time in the real RAM model of computation.

Theorem 1 *For every ordered tree T with $n \geq 2$ leaves, without vertices of degree two, and with edge lengths for edges connecting non-leaves, we can find a set S of n sites in $O(n)$ time such that the farthest-point Voronoi diagram of S is a realization of T where every bounded edge in $\mathrm{fVor}(S)$ has a prescribed length.*

4 Other Distance Functions

Voronoi diagrams and farthest-point Voronoi diagrams can naturally be generalized to a wider class of distance functions defined as follows: a distance function d is specified by giving its *unit circle* C_d, i.e., all those points considered to have distance one from the origin. We assume throughout that d is convex and symmetric, i.e., C_d is a closed curve that bounds a convex shape that has 2-fold rotational symmetry about the origin.

To measure distances, we use *homothets* of C_d, i.e., scaled and translated copies. We call such a homothet a *d-disc* and say that it is *centered at* p if the origin was translated to p. Given a set S of sites, let the *full d-disc* $F^d_S(p)$ be the smallest d-disc centered at p that encloses all sites of S. The *d-farthest-point Voronoi diagram* of a set S of sites, denoted by $\mathrm{fVor}_d(S)$, is defined as before by letting p be a vertex (resp. interior point of an edge) if and only if $F^d_S(p)$ contains three (resp. two) sites.[2]

We briefly argue that this indeed expresses "farthest" correctly. For two points p and q, the distance $d(p, q)$ (with respect to the distance function defined by C_d) is defined to be the smallest scaling factor at which a d-disc centered at p touches q. Since d is symmetric, we have $d(p, q) = d(q, p)$. In particular, a site $s \in S$ is farthest from the point p if s is on the boundary of $F^d_S(p)$. If p is a point on an edge of $\mathrm{fVor}_d(S)$, then by definition there are two sites s, s' on $F^d_S(p)$. Thus p is equidistant from s, s' and all other sites are no farther. Hence any edge of $\mathrm{fVor}_d(S)$ bounds a region where all points have the same farthest point. See Figure 4.

4.1 Smooth Strictly Convex Symmetric Distances

We call a distance function d *strictly convex* if the boundary of C_d contains no line segments, and *smooth* if every point on the boundary of C_d has a unique tangent. We now show that we can realize arbitrary ordered trees as d-farthest-point Voronoi diagram for any smooth and strictly convex symmetric distance function d.

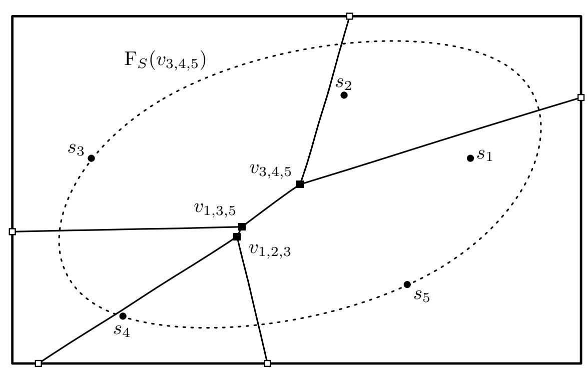

Figure 4: A d-farthest-point Voronoi diagram.

The approach is the same as for the Euclidean case, with the only change that we use C_d, rather than geometric circles, to define arcs to place sites on. Thus, for a tree T_1 with a single interior node v with ℓ incident leaves, place ℓ sites on the unit circle C_d. The origin becomes the Voronoi vertex that we identify with v.

[2]For non-symmetric convex distances, the full d-disc is a mirrored homothet of C_d and the correspondence to vertices and edges of the diagram no longer holds [9].

To create sites for a tree T_{k+1} with $k+1$ interior nodes, find one node v that is adjacent to only one other interior node u, and remove all incident leaves of v. Recursively find sites for the resulting tree T_k. Find the unbounded edge r from u on which the symbolic endpoint for v resides, and pick an arbitrary point p on it. Find the full d-disc $\mathrm{F}_S^d(p)$; this contains the two sites s, s' whose farthest regions meet at edge r on their boundaries. Turn p into a vertex of the d-farthest-point Voronoi diagram by placing sites at the shorter arc of $\mathrm{F}_S^d(p)$, placing $\ell - 1$ sites if v was incident to ℓ leaves.

It remains to argue that this is correct, i.e., that all newly placed sites are safe and proper. In a nutshell, this holds because they are strictly inside $\mathrm{F}_S^d(u)$ and strictly outside $\mathrm{CH}(S)$. We give a proof in the appendix.

Theorem 2 *Let d be a smooth and strictly convex symmetric distance function. For every ordered tree T with $n \geq 2$ leaves, without vertices of degree two, and with edge lengths for edges connecting non-leaves, we can find a set S of n sites in $O(n)$ time such that the d-farthest-point Voronoi diagram of S is a realization of T where every bounded edge has its prescribed length.*

4.2 Polygonal Convex Symmetric Distances

We now illustrate some of the challenges that arise when our distance function is not smooth or not strictly convex. Unlike for strictly convex distances, the *d-bisector* of two sites s and s' (i.e., the set of all points that are equidistant from s and s' with respect to d) is not necessarily homeomorphic to a line, and indeed, may be a 2-dimensional region. Ma [9] shows that this occurs precisely when the line segment ss' is parallel to a line segment on the boundary of the unit circle C_d that defined d. This limits our ability to realize ordered trees as d-farthest-point Voronoi diagrams when d is *polygonal*, i.e., C_d is a k-sided convex polygon.

Theorem 3 *Let d be a convex distance function defined by a polygon with k edges and let T be a tree with more than k leaves. There is no set of sites S such that the d-farthest-point Voronoi diagram of S realizes T.*

Proof. For every edge e of the unit circle C_d, at most one site can be extreme in the direction normal to e. More precisely, for any half-plane $h \supset S$ whose bounding line ℓ is parallel to e, there is at most one site on ℓ—otherwise $\mathrm{fVor}_d(S)$ is not a tree. So if $\mathrm{fVor}_d(S)$ is a tree, then at most k sites in S have nonempty cells, hence the tree has at most k leaves. Therefore, we cannot realize trees with more than k leaves. □

For example, for the L_1-distance and the L_∞-distance, the unit circle C_d is a 4-sided polygon, so no tree with more than four leaves can be realized as farthest-point Voronoi diagrams under these distances.

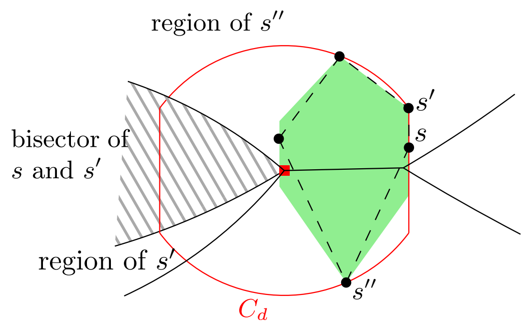

Figure 5: If some portion of C_d (red) is a line segment, two sites on that line segment (e.g., s, s') can have a two-dimensional bisector (grey region). The generalized convex hull $\mathcal{H}(S)$ (green) may strictly include the convex hull (dashed). Here, the site s is a vertex of the (ordinary) convex hull but s is not proper: removing s leaves the generalized convex hull unchanged.

A second problem with distance functions that are not strictly convex is that not all extreme points of the convex hull are proper; for example point s in Figure 5 is an extreme point of $\mathrm{CH}(S)$ but any point p for which s is farthest also has s' as farthest point.

However, we can prove a similar relationship. Let $\mathcal{H}(S)$ be the intersection of all d-discs that contain S. We refer to $\mathcal{H}(S)$ as the *generalized convex hull* of S. We call a site s an *extreme point* of $\mathcal{H}(S)$ if $\mathcal{H}(S) \neq \mathcal{H}(S \setminus \{s\})$. We give in the appendix the following characterization:

Lemma 4 *A site s in S is proper if and only if s is an extreme point of the generalized convex hull $\mathcal{H}(S)$.*

We may attempt to follow the steps of the algorithm from the Euclidean setting, in the hope of always finding proper sites. We now show that this can fail. As before define v, u, r, s, s' in the expansion step. Presume we are in a situation where $\mathrm{F}_S^d(u)$ contains s, s' on adjacent straight-line edges. Then the generalized hull $\mathcal{H}(S)$ coincides with $\mathrm{F}_S^d(u)$ on the stretch between s and s'. Thus, the region where we placed sites for strictly convex distances is empty, giving no suitable, safe, proper candidates. Put differently, we cannot longer realize ordered trees in the carefree online fashion we use for the Euclidean distance. Rather, we need to know the ordered tree in advance and we need to decide *a priori* which site will occupy which edge of C_d. We conjecture that with a judicious choice, we can realize every tree with at most k leaves if C_d is a k-sided polygon, but this remains an open problem. Without giving details, we note that all ordered trees *can* be realized by any convex symmetric distance function for which C_d is strictly convex and smooth in at least one region, by placing all initial sites and later additions only within that part of C_d.

5 Geometric Trees

In this section, we study how to test whether a specified geometric tree is a farthest-point Voronoi diagram in the Euclidean metric. We are given a tree with a fixed drawing in the plane, with the leaves at infinity. Reinterpreting this, we are given a subdivision of the plane into cells, and we ask whether there exists a set of sites whose farthest-point Voronoi diagram comprises these cells. An affirmative answer is possible only if every cell of the subdivision is convex and unbounded.

Our approach generalizes to arbitrary dimension k, so assume that we are given a convex subdivision Z of $\mathbb{R}^k$, where each cell in Z is a convex, unbounded polyhedron. We wish to determine whether Z is the farthest-point Voronoi diagram of some set S of sites. Each cell in Z has some number of $(k-1)$-dimensional facets (e.g., edges if $k=2$), and we assume that for each such facet f we know a unit normal vector n_f. Thus, for each facet f, its affine hull has the form $\{p : \langle n_f, p\rangle = \alpha_f\}$, where α_f is a suitable scalar. Let $f_{\sigma\tau}$ denote a facet whose incident cells are σ and τ, where n_f is directed from τ into σ and thus σ is the cell whose interior points have a positive signed distance from $f_{\sigma\tau}$ (i.e., $\langle n_f, p\rangle \geq \alpha_f$ for all points $p \in \sigma$.)

Suppose Z can be realized as farthest-point Voronoi diagram. In this realization each cell σ is assigned a site $\rho(\sigma)$ such that the points in σ are exactly those points for which $\rho(\sigma)$ is the farthest site. We will describe any (putative) realization as such a function $\rho(\sigma)$.

The following result holds for realizations of farthest-point Voronoi diagrams in arbitrary dimension (and also for ordinary Voronoi diagrams [5]).

Lemma 5 (bisector condition) *Let ρ be a realization of Z. For every facet $f_{\sigma\tau}$ in Z, the affine hull of $f_{\sigma\tau}$ must be the bisector of $\rho(\sigma)$ and $\rho(\tau)$.*

Hence, given $\rho(\sigma)$ we can compute $\rho(\tau)$ by reflecting $\rho(\sigma)$ about f, i.e., $\rho(\tau) = \rho(\sigma) - 2(\langle n_f, \rho(\sigma)\rangle - \alpha_f)n_f$. As this is an affine equation in $\rho(\sigma)$, it can be expressed in the matrix form

$$\begin{bmatrix} \rho(\tau) \\ 1 \end{bmatrix} = R_{\sigma\tau} \begin{bmatrix} \rho(\sigma) \\ 1 \end{bmatrix}$$

where $R_{\sigma\tau}$ is a $(k+1)\times(k+1)$ matrix determined solely by the normal vector and scalar of the face $f_{\sigma\tau}$. Thus we have a system of $k+1$ equations for each facet of Z. Let $\bar{\tau}$ denote the vector $[\rho(\tau) \;\; 1]^{\mathrm{T}}$, so the equation becomes $\bar{\tau} = R_{\sigma\tau}\bar{\sigma}$.

We need a second condition. In the ordinary Voronoi diagram, a site must lie inside the cell of points for which it is the nearest site. For the farthest-point Voronoi diagram, we need a condition that is essentially the inverse.

Lemma 6 (outside condition) *Let ρ be a realization of a subdivision Z. For every facet f incident to a cell σ, the affine hull H of f has the cell σ on one side and the site $\rho(\sigma)$ on the other.*

Proof. Say the facet is $f = f_{\sigma\tau}$. According to the bisector condition, $\rho(\sigma)$ and $\rho(\tau)$ are on opposite sides of H, and every point on the same side of H as $\rho(\sigma)$ is closer to $\rho(\sigma)$ than it is to $\rho(\tau)$. No point $p \in \sigma$ can lie on the same side of H as $\rho(\sigma)$, as p's farthest site cannot be $\rho(\sigma)$. □

We express the outside condition as the two inequalities

$$\langle n_{\sigma\tau}, \rho(\sigma)\rangle \leq \alpha_{\sigma\tau} \leq \langle n_{\sigma\tau}, \rho(\tau)\rangle,$$

where, as before, $n_{\sigma\tau}$ is a unit vector normal to $f_{\sigma\tau}$ such that $\langle n_{\sigma\tau}, p\rangle \geq \alpha_{\sigma\tau}$ for all points $p \in \sigma$ and $\langle n_{\sigma\tau}, p\rangle \leq \alpha_{\sigma\tau}$ for all points $p \in \tau$. Crucial to our testing routine is the following.

Theorem 7 *Let Z be a convex subdivision of $\mathbb{R}^k$. Let $S = \rho(\cdot)$ be an assignment of sites to cells in Z. Then Z is the farthest-point Voronoi diagram of S if and only if the bisector condition and the outside condition holds for every facet of Z.*

Proof. Necessity has been shown already. Suppose for the sake of contradiction that the two conditions hold, yet Z is not the farthest-point Voronoi diagram of S. Then there exists some cell σ of Z containing an interior point p for which the farthest site in S is not $\rho(\sigma)$ but instead some other site $\rho(\tau)$ assigned to a cell τ.

Shoot a ray from the interior of τ toward p, and let $f_{\tau\omega}$ be the first facet (breaking ties arbitrarily) of τ that the ray strikes. The ray strikes $f_{\tau\omega}$ before reaching p, as p is in the interior of a cell other than τ; therefore, p is on ω's side of the affine hull of $f_{\tau\omega}$. By the outside condition therefore p is *not* on $\rho(\omega)$'s side of the affine hull of $f_{\tau\omega}$. As $f_{\tau\omega}$ bisects $\rho(\omega)$ and $\rho(\tau)$ by the bisector condition, therefore p is closer to $\rho(\tau)$ than to $\rho(\omega)$, contradicting the fact that $\rho(\tau)$ is the site in S that is farthest from p. The result follows. □

Theorem 7 implies that we can answer the question by finding a set S of sites that satisfy all the bisector conditions and outside conditions—one of the former and two of the latter for each facet of Z—or by showing that no such set of sites exists. As the bisector conditions are linear equations and the outside conditions are linear inequalities, the question reduces to finding a feasible point of a linear program.

For efficiency, we recommend reducing the linear program to k variables prior to solution by performing substitutions of the bisector conditions. We achieve this with a *propagation* procedure that exploits the dual graph of the convex subdivision Z, as Biedl et al. [5] do for the ordinary Voronoi diagram. Form the dual graph G of Z: G's vertices correspond to the k-cells of

Z and G's edges correspond to Z's facets. Choose a distinguished k-cell σ in Z (hence a distinguished node in the graph). The variables in our system are the coordinates of the putative site $\rho(\sigma)$, hence the first k entries of vector $\bar{\sigma}$. Perform a depth-first search of G, during which we express the coordinates of every other site as a linear combination of $\bar{\sigma}$'s coordinates by composing reflections of the form $\bar{\tau} = R_{\omega\tau}\bar{\omega}$. Composing these reflections is simply matrix multiplication; thus we obtain a linear relationship of the form $\bar{\tau} = R'_{\sigma\tau}\bar{\sigma}$ for every cell τ, even those that do not share a facet with σ. ($R'_{\sigma\tau} = R_{\sigma\tau}$ if (σ, τ) is an edge of G.)

Next, consider the edges of G that the depth-first search did not traverse. Each such edge (ω, τ) corresponds to a facet of Z that introduces an additional reflection equation of the form $\bar{\omega} = R_{\tau\omega}\bar{\tau}$, which hence becomes another linear equality constraint imposed on $\bar{\sigma}$: $R'_{\sigma\omega}\bar{\sigma} = R_{\tau\omega}R'_{\sigma\tau}\bar{\sigma}$. However, these constraints are often redundant or trivial (i.e., $\bar{\sigma} = \bar{\sigma}$). We can stack these linear equations ($k+1$ equations per untraversed edge) in the form of a matrix equation $M\bar{\sigma} = b$, where M has $k+1$ columns and $O(mk)$ rows, and m is the number of facets in Z. This linear system hence defines an affine subspace Λ of vectors $\bar{\sigma}$ that are compatible with the bisector condition. Typically Λ is a single point or empty, but it could have dimension as high as k.

The outside condition imposes another system of $O(mk)$ linear inequalities, two per facet. If Λ is a single point, it is now a simple matter to check whether it satisfies all these inequalities. If Λ is a larger subspace, we restrict the inequalities to the subspace Λ and solve the consequent linear program. Any feasible point can be used for $\bar{\sigma}$ (hence gives the site $\rho(v)$), and we can compute the other sites by applying the reflection equations. The solution space may have dimension up to k, as Figure 6 illustrates. If $\Lambda = \emptyset$ or the linear program is infeasible, Z is not a farthest-point Voronoi diagram of any set of sites.

Suppose Z has n cells and m facets in k dimensions. It takes $O(nk^3)$ time to compute the propagation matrices $R'_{\sigma\tau}$ (accounting for fewer than n multiplications of $(k+1)\times(k+1)$ matrices); $O(mk^3)$ time to compute the remaining equations and inequalities due to the bisector and outside conditions; and $O(f(k)(n+m))$ time to solve the linear program where $f(\cdot)$ is a function (typically exponential) [10]. As the size of the input subdivision Z is $\Omega(m+n)$, the total running time is linear in the input size if the dimension k is a small constant.

Theorem 8 *Given a convex subdivision Z of $\mathbb{R}^k$, where k is a small constant, we can test in linear time whether there exists a set of sites whose farthest-point Voronoi diagram is Z.*

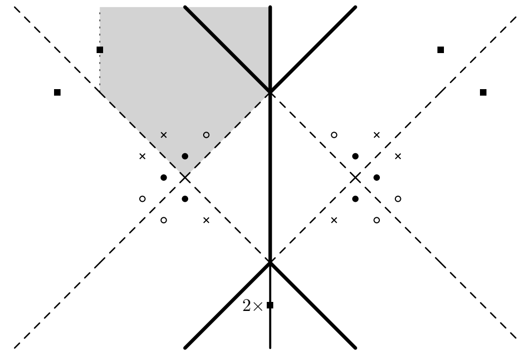

Figure 6: A farthest-point Voronoi diagram (thick edges) and three sets of sites (discs, circles, crosses) that realize it. Once one site is fixed in the open gray cell $\mathcal{G}$, the others follow by reflection at the bisectors (thick or dashed). Sites on the boundary of $\mathcal{G}$ (e.g., the square) yield sites that coincide, and sites outside $\mathcal{G}$ generate sites that violate the outside condition.

References

[1] O. Aichholzer, T. Biedl, T. Hackl, M. Held, S. Huber, P. Palfrader, and B. Vogtenhuber. Representing directed trees as straight skeletons. In *Graph Drawing and Network Visualization (GD '15)*, pages 335–347, 2015.

[2] O. Aichholzer, H. Cheng, S. L. Devadoss, T. Hackl, S. Huber, B. Li, and A. Risteski. What makes a tree a straight skeleton? In *Canadian Conference on Computational Geometry (CCCG '12)*, pages 253–258, 2012.

[3] G. Aloupis, H. Pérez-Rosés, G. Pineda-Villavicencio, P. Taslakian, and D. Trinchet-Almaguer. Fitting Voronoi diagrams to planar tesselations. In *Intl. Workshop on Combinatorial Algorithms (IWOCA 2013)*, pages 349–361, 2013.

[4] P. Ash and E. Bolker. Recognizing Dirichlet tesselations. *Geometriae Dedicata*, 19:175–206, 1985.

[5] T. Biedl, M. Held, and S. Huber. Recognizing straight skeletons and Voronoi diagrams and reconstructing their input. In *10th International Symposium on Voronoi Diagrams in Science and Engineering (ISVD 2013)*, pages 37–46, 2013.

[6] D. Hartvigsen. Recognizing Voronoi diagrams with linear programming. *ORSA J. Comput.*, 4:369–374, 1992.

[7] G. Liotta and H. Meijer. Voronoi drawings of trees. In *Graph Drawing (GD '99)*, pages 369–378, 1999.

[8] G. Liotta and H. Meijer. Voronoi drawings of trees. *Comput. Geom.*, 24(3):147–178, 2003.

[9] L. Ma. *Bisectors and Voronoi Diagrams for Convex Distance Functions.* PhD thesis, FernUniversität Hagen, 2000.

[10] N. Megiddo. Linear programming in linear time when the dimension is fixed. *J. ACM*, 31(1):114–127, 1984.

A Smooth Strictly-Convex Distance Functions

Recall that the distance function d is given by specifying its *unit circle* C_d, a *d-disc* is a homothet of C_d, and the *radius* of a d-disc D is the scaling factor used to obtain D from C_d. In this section, we show in detail that if C_d is strictly convex and smooth, then our algorithm to find sites whose farthest-point Voronoi diagram realizes a given ordered tree T works correctly. There are two things that must be shown: every added site s is *d-proper* (there exists a point p for which s is the unique farthest site) and *d-safe* (all previously placed sites remain d-proper).

A.1 Proper Sites

Recall that an *extreme point* of the convex hull $\mathrm{CH}(S)$ is a site $s \in S$ such that $\mathrm{CH}(S \setminus \{s\})$ is a strict subset of $\mathrm{CH}(S)$. Equivalently, a site $s \in S$ is an extreme point of S if there exists a half-space ℓ that has all points in $S \setminus \{s\}$ in its interior and s in its exterior.

Theorem 9 *Let S be a set of sites in the plane and let d be a smooth strictly convex distance function.*

1. *A site s is d-proper if and only if s is an extreme point of the convex hull of S.*
2. *The regions of two sites s_i and s_j share an unbounded edge if and only if s_i and s_j are consecutive extreme points of the convex hull of S.*
3. *The d-proper sites appear in the same order along the convex hull of S as their corresponding regions in the d-farthest-point Voronoi diagram.*

Proof. To show the first claim, suppose the site s is d-proper. Then there is a point p such that $\mathrm{F}^d_S(p)$ has only the site s on its boundary. Since C_d is convex, the convex hull $\mathrm{CH}(S)$ is contained in $\mathrm{F}^d_S(p)$. Since C_d is strictly convex, $\mathrm{CH}(S)$ intersects $\mathrm{F}^d_S(p)$ only in point s. Hence, $\mathrm{CH}(S\setminus\{s\})$ is strictly inside $\mathrm{F}^d_S(p)$, which proves that $\mathrm{CH}(S \setminus \{s\}) \subset \mathrm{CH}(S)$ and, thus, the site s is an extreme point of the convex hull $\mathrm{CH}(S)$.

Conversely, suppose s is an extreme point of $\mathrm{CH}(S)$, say half-space ℓ separates s from the rest of S. Since C_d is smooth, there exist two points on C_d whose tangent has the same slope as the affine hull of ℓ. By scaling C_d sufficiently much, we can hence find a homothet D of C_d that in the vicinity of one of these points is arbitrarily close to ℓ. Hence D contains $S \setminus \{s\}$ and not s. Scaling D while keeping its center then yields a d-disc with only s on its boundary, proving that the region of s is non-empty.

The proof of (2) and (3) is very similar to part (1) after observing that (s_i, s_j) is an edge of the convex hull if and only if there exists a half-space ℓ that contains all points in $S \setminus \{s_i, s_j\}$ in its interior and s_i, s_j in its exterior. With this we can find an unbounded region of points whose farthest site is either s_i or s_j, and therefore there must be an unbounded edge separating their two regions. □

As we will see below, we always choose the next site(s) to be outside the convex hull of the current sites. As such, all sites that we choose will be d-proper.

A.2 Properties of Homothets

Before proving safety, we need some basic observations about homothets of a strictly convex smooth C_d.

Theorem 10 (Ma [9]) *Let D and D' be two different homothets of a compact convex set C_d. Then the boundaries of D and D' intersect in at most two points, or in a point and a line segment, or in two line segments.*

Corollary 11 *Let D and D' be two different homothets of a strictly convex smooth compact set C_d. Then the boundaries of D and D' intersect at most two points.*

Proof. The claim follows from Theorem 10, since the boundary of a homothet of a strictly convex compact set does not contain any line segments, by definition. □

We say that two curves C, C' *truly intersect* at some point p if they have p in common, and any sufficiently small circle centered at p intersects the curves in four points and in order C, C', C, C'.

Lemma 12 *Let D and D' be two different homothets of a strictly convex smooth compact set C_d. If the boundaries of D and D' intersect in two points a, b, then they truly intersect at both a and b.*

Proof. We consider the situation near a. Since D and D' are smooth, there are unique tangents t_a and t'_a at a for D and for D', respectively. We argue that these tangents have different slopes.

Since C_d is strictly convex, the slope of the tangent determines the point on C_d uniquely, up to reflection through the center-point, and the line from this point to the center-point has the same slope regardless of how we scale or translate C_d. Thus, the line from a to the center-point p of D has the same slope as the line from a to the center-point p' of D', so p, a, p' are all on one line.

Repeating the argument at b, we see that p, b, p' (and therefore also a) are all on one line. But then D and D' must have the same scale-factor (else they could not both contain both a and b), and therefore the same center-point, and so are the same homothet. Contradiction, so t_a and t'_a have different slopes. Since D and D' are smooth, their boundary locally follows the lines along t_a and t'_a, which means that they truly intersect at a. □

Finally we need a rather technical observation, which will be crucial for defining the "lunes" which are used for placing sites safely.

Lemma 13 (Inside-Outside Lemma) *Let a and b two points in the plane and let h and $\bar{h}$ be the half-planes bounded by the line through a and b. Consider two d-discs D and D' such that*

(a) the centers of D and D' both lie in h,

(b) the radius of D' is larger than the radius of D, and

(c) the boundaries of D and D' intersect at a and b.

Then we have the following.

(1) Within the half-plane h, the d-disc D' contains D, i.e., $h \cap D \subset h \cap D'$.

(2) Within the half-plane $\bar{h}$, the d-disc D contains D', i.e., $\bar{h} \cap D' \subset \bar{h} \cap D$.

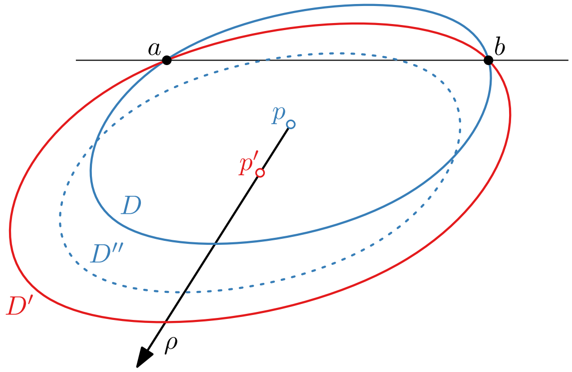

Figure 7: Two d-discs D (blue) and D' (red) that have their centers p and p' on the same side as the line through their two intersection points a and b. The ray ρ from p through p' first hits D, then ρ hits a copy D'' of D centered at p' (dotted, blue), and finally ρ hits D'.

Proof. Let p be the center of D and p' the center of D'. Consider the ray ρ that shoots from p through p'. We argue that ρ hits D strictly before D'.

As illustrated in Figure 7, we place a copy D'' of D centered at p'. The ray ρ hits D before D'', since D'' is a copy of D translated from p to p'. Furthermore, the ray ρ hits D'' strictly before D', since D' is a strictly larger copy of D'' with the same center. This means that the ray ρ hits the boundary of D strictly before the boundary of D'. Since D and D' are strictly convex and homothetic, the boundaries of D and D' cannot have any intersection other than a and b. Therefore, within the half-space h, the boundary of D lies in the interior of D', i.e., $h \cap D \subset h \cap D'$. This proves (1).

To show (2), observe that since the boundaries of D and D' intersect in two points, at both points we have true intersections. Due to (1), we *enter* D as we traverse the boundary of D' from h to $\bar{h}$ through a (or through b). Since the boundaries of D and D' intersect only at a and b, we know that, within $\bar{h}$, the boundary of D' lies in the interior of D, i.e., $\bar{h} \cap D' \subset \bar{h} \cap D$. □

A.3 Lunes and Safe Sites

Let us assume that the sites are numbered $s_1, s_2, \ldots, s_n$ in an arbitrary manner. Let $v_{i,j,k}$ be the point equidistant to sites s_i, s_j, and s_k; and let $e_{i,j}$ be the edge (if any) on the bisector of sites s_i and s_j. Suppose p is a point along an unbounded edge $e_{i,j}$ defined by the sites s_i and s_j, and we want to place a new site s on the d-arc $\mathrm{A}_d(p, s_i, s_j)$ to create a new vertex at some point p. Define the d-*lune* $\mathrm{Lune}_d(s_i, s_j)$ to be the union of all d-arcs $\mathrm{A}_d(p, s_i, s_j)$ such that p is an interior point of ray r. Figure 8 depicts an example of a d-lune.

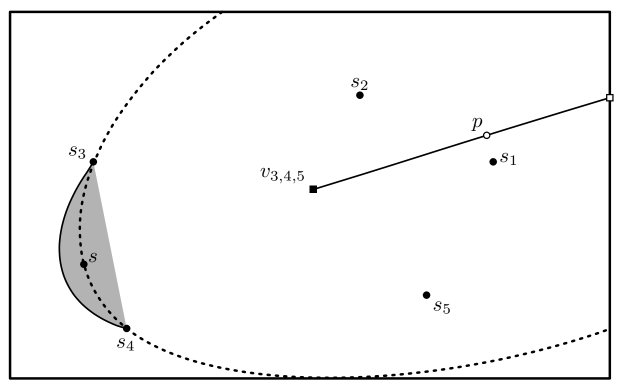

Figure 8: The d-lune $\mathrm{Lune}_d(s_3, s_4)$ for the sites from Figure 4 together with its defining edge $e_{3,4}$. A new site s in this d-lune creates a new vertex at p along $e_{3,4}$, where p is the center of the d-disc through s_3, s_4, and s.

Lemma 14 *For any two consecutive vertices s_i, s_j on $\mathrm{CH}(S)$, if $v_{i,j,k}$ is the finite end of edge $e_{i,j}$, then any point in $\mathrm{Lune}_d(s_i, s_j)$ belongs to $\mathrm{F}^d_S(v_{i,j,k}) \setminus \mathrm{CH}(S)$.*

Proof. Consider $\mathrm{F}^d_S(p)$ for some point p on $e_{i,j}$. By definition of a full circle it contains all sites in S, so $\mathrm{CH}(S) \subset \mathrm{F}^d_S(p)$ since C_d is convex. Therefore $A(p, s_i, s_j)$ is outside $\mathrm{CH}(S)$. On the other hand, both p and $v_{i,j,k}$ are within one half-plane h defined by the line through s_i, s_j (since $e_{i,j}$ consists of those points for which these are the farthest sites). By the Inside-Outside lemma therefore $A(p, s_i, s_j)$ (which is outside h) therefore is within $\mathrm{F}^d_S(v_{i,j,k}) \cap \bar{h}$. □

So as promised previously, all newly placed sites are outside the convex hull of preexisting sites, and so are proper. Now we are ready to prove safety.

Lemma 15 (Safety Lemma) *For any two consecutive vertices s_i, s_j on* $\mathrm{CH}(S)$*, every new site in* $\mathrm{Lune}_d(s_i, s_j)$ *is safe.*

Proof. Let s be be a new site for S that is contained in $\mathrm{Lune}_d(s_i, s_j)$. Let $e_{i,j}$ be the unbounded edge where the regions of s_i and s_j meet, and let $v_{i,j,k}$ be the vertex where $e_{i,j}$ ends. By the definition of $\mathrm{Lune}_d(s_i, s_j)$, the new site s is contained in the full d-disc $\mathrm{F}^d_S(v_{i,j,k})$ that passes through s_i and s_j. Thus, s is safe for $v_{i,j,k}$.

Consider a vertex $v_{i,k,l}$ that is connected to $v_{i,j,k}$ by the edge $e_{i,k}$. We argue that $\mathrm{Lune}_d(s_i, s_j)$—and, therefore, the new site s—is contained in $\mathrm{F}^d_S(p)$ for any point $p \in e_{i,k}$, i.e., the new site s is safe for $e_{i,k}$ and $v_{i,k,l}$.

Let h_s be the half-plane containing s that is bounded by the line through s_i and s_k. We apply Lemma 13 in two ways, depending on whether p lies in h_s or not.

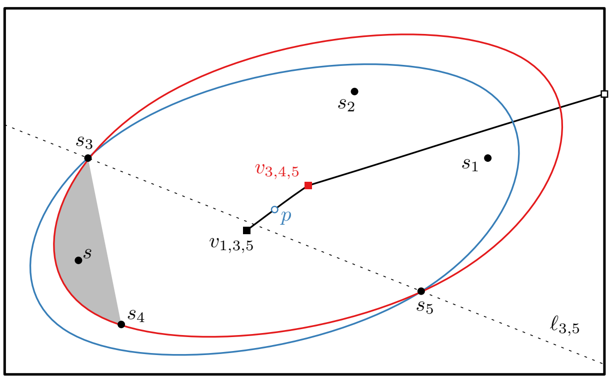

Figure 9: An example for the case $p \notin h_s$ from the proof of Lemma 15 with $i = 3$, $j = 4$, $k = 5$, and $l = 1$.

Suppose $p \notin h_s$, as illustrated in Figure 9. We approach s_i and s_k when we walk from $v_{i,j,k}$ along $e_{i,k}$ towards $v_{i,k,l}$. Therefore, $\mathrm{F}^d_S(v_{i,j,k})$ is larger than $\mathrm{F}^d_S(p)$. Since $p, v_{i,j,k} \notin h_s$, Lemma 13 implies $h_s \cap \mathrm{F}^d_S(v_{i,k,l}) \subset h_s \cap \mathrm{F}^d_S(p)$. We know $s \in \mathrm{Lune}_d(s_i, s_j) = h_s \cap \mathrm{F}^d_S(v_{i,k,l})$. Therefore, $s \in h_s \cap \mathrm{F}^d_S(p)$, and, thus, s is safe for p.

Suppose $p \in h_s$, as illustrated in Figure 10. Then there is a point w along $e_{i,k}$ that intersects $\ell_{i,k}$, since $v_{i,j,k} \notin h_s$. We move away from s_i and s_k when we walk from w along $e_{i,k}$ to $v_{i,k,l}$. Therefore, $\mathrm{F}^d_S(p)$ is larger than $\mathrm{F}^d_S(w)$. Since $p, w \in h_s$, Lemma 13 implies $h_s \cap \mathrm{F}^d_S(w) \subset h_s \cap \mathrm{F}^d_S(p)$. We know from the previous case, when $p \notin h_s$, that $s \in h_s \cap \mathrm{F}^d_S(w)$. Therefore, $s \in h_s \cap \mathrm{F}^d_S(p)$ and, thus, the new site s is safe for p.

In summary, if $s \in \mathrm{Lune}_d(s_i, s_j)$ is safe for $v_{i,j,k}$ then s is safe for all edges incident to $v_{i,j,k}$, except for the unbounded edge $e_{i,j}$. We can repeat the above argument for all neighbors of $v_{i,j,k}$ and their neighbors and so forth. In this fashion, the safety of s propagates to all vertices and all bounded edges of the d-farthest-point Voronoi diagram of S.[3] Therefore s is safe for S. □

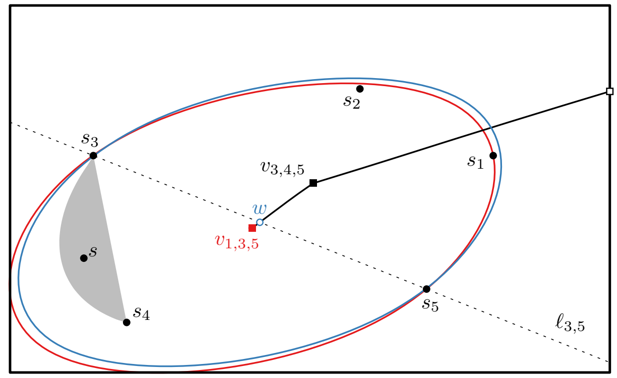

Figure 10: An example for the case $p \in h_s$ from the proof of Lemma 15 with $i = 3$, $j = 4$, $k = 5$, and $l = 1$.

B Polygonal Distance Functions: Proof of Lemma 4

Proof. Suppose s is a proper site in S. Then there is a point p such that $\mathrm{F}^d_S(p)$ has only the site s on its boundary. All other sites of S are in the interior of $\mathrm{F}^d_S(p)$ by definition of full disc. Scaling $\mathrm{F}^d_S(p)$ down while staying centered at p gives another homothet D of C_d; note that $D \subset \mathrm{F}^d_S(p)$ since d is convex. If we shrink little enough then D hence contains all of $S \setminus \{s\}$, but it does not contain s. Therefore, $\mathcal{H}(S \setminus \{s\}) \subseteq D$ does not contain s. By definition, s is an extreme point of $\mathcal{H}(S)$.

Conversely, suppose s is an extreme point of $\mathcal{H}(S)$. That means there is a homothet D of C_d that contains $S \setminus \{s\}$ and that does not contain s. Let p be the center of D. Suppose we grow D until we arrive at a d-disc D' centered at p with s on the boundary. We have $D \subset D'$, since both D and D' are convex and symmetric to p. Hence, D' is a d-disc centered at p that contains S and has only the site s on its boundary. This means s is the only d-farthest point from p, i.e. s is a proper site. □

[3] In fact, the safety of s extends to all unbounded edges other than $e_{i,j}$ in the diagram for S, as well.