Piecewise-Linear Farthest-Site Voronoi Diagrams

Franz Aurenhammer
Institute for Theoretical Computer Science, TU Graz, Austria

Evanthia Papadopoulou
Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland

Martin Suderland
Faculty of Informatics, Università della Svizzera italiana, Lugano, Switzerland

Abstract

Voronoi diagrams induced by distance functions whose unit balls are convex polyhedra are piecewise-linear structures. Nevertheless, analyzing their combinatorial and algorithmic properties in dimensions three and higher is an intriguing problem. The situation turns easier when the farthest-site variants of such Voronoi diagrams are considered, where each site gets assigned the region of all points in space farthest from (rather than closest to) it.

We give asymptotically tight upper and lower worst-case bounds on the combinatorial size of farthest-site Voronoi diagrams for convex polyhedral distance functions in general dimensions, and propose an optimal construction algorithm. Our approach is uniform in the sense that (1) it can be extended from point sites to sites that are convex polyhedra, (2) it covers the case where the distance function is additively and/or multiplicatively weighted, and (3) it allows an anisotropic scenario where each site gets allotted its particular convex distance polytope.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Voronoi diagram, farthest-site, polyhedral distance, polyhedral sites, general dimensions

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2021.30

Funding The first author was supported by Projects I 1836-N15 and I 5270-N, Austria Science Fund (FWF). The last two authors were supported in part by Projects 200021E_154387 and 200021E_201356, Swiss National Science Foundation (SNF).

1 Introduction

The *Voronoi diagram* of a set of n geometric objects, called *sites*, is a well-known space partitioning structure with numerous applications in diverse fields of science. In its *closest-site* variant, this diagram partitions the underlying space into maximal regions such that all points within one region have the same closest site. The Euclidean Voronoi diagram of point sites in $\mathbb{R}^d$ is well studied; see e.g. [4, 7, 11, 15]. It is a piecewise-linear cell complex of worst-case complexity $\Theta(n^{\lceil \frac{d}{2} \rceil})$ and can be constructed in $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$ time. There are many ways to modify this standard diagram, for example, by using different distance measures, by considering sites of more general shape, or by assigning individual weights to them. Adapting to practical needs, such generalizations (among several others) have been studied extensively, and many satisfactory results are available nowadays [4, 18].

For most of these generalized Voronoi diagrams, the partition of space they define is not piecewise linear any more, but is rather composed of curved geometric objects of various dimensions and shapes. This complicates their structural analysis as well as their computational construction, especially in dimensions higher than two (where results are becoming comparatively sparse). For instance, already in three-dimensional space $\mathbb{R}^3$, the algebraic description of the edges and facets of the Euclidean Voronoi diagram of straight lines becomes exceedingly complicated [12]. What is more, the combinatorial complexity of

32nd International Symposium on Algorithms and Computation (ISAAC 2021).
Editors: Hee-Kap Ahn and Kunihiko Sadakane; Article No. 30; pp. 30:1–30:11
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

this diagram is a major open problem in computational geometry [17]. There is a gap of an order of magnitude between the $\Omega(n^2)$ lower bound [1] and the only known upper bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$ [21].

Certain types of Voronoi diagrams retain their piecewise linear structure, however. For example, the so-called power diagram [2] has this property. Another prominent class, and the one of interest in the present note, is induced by (convex) *polyhedral distance functions*. Intuitively speaking, the distance from a point x to a site s is now measured as the extent at which a given convex polyhedron $\mathcal{P}$, which being centered at x, has to expand till it starts touching s.

Several authors succeeded in deriving strong bounds on the combinatorial complexity of such Voronoi diagrams. If the n sites are points and the distance polytope $\mathcal{P}$ is a simplex or a cube – the latter just giving the L_∞ distance – then this complexity in $\mathbb{R}^d$ is $\Theta(n^{\lceil \frac{d}{2} \rceil})$, see [6]. (The dimension d is considered a constant throughout this paper.) In $\mathbb{R}^3$, the same bound still applies when any constant-sized convex polytope is chosen for $\mathcal{P}$ [13]. For the sites being n straight lines in $\mathbb{R}^3$, with $\mathcal{P}$ defined as before, near-quadratic bounds of $\Omega(n^2\alpha(n))$ and $O(n^2\alpha(n)\log(n))$ can be obtained [9]. Here $\alpha(n)$ is the extremely slowly growing inverse Ackermann function. If we consider as sites disjoint convex polyhedra with n faces in total, then the complexity is $O(n^{2+\epsilon})$, as has been shown in [16]; this sharpens to $O(n^2\alpha(n)\log(n))$ if all the sites are line segments.

Though Voronoi diagrams for convex polyhedral distance functions – in comparison to Euclidean Voronoi diagrams – thus proved easier to deal with concerning their combinatorial aspects, this does not seem to carry over to their algorithmic aspects. In fact, the papers cited above do not provide algorithms for computing such diagrams, and we are not aware of any construction algorithm particular to them.

As we shall show in this note, the situation changes if the so-called *farthest-site* variant of the diagram is considered (rather than the closest-site variant as above). The farthest-site Voronoi diagram is a partition of the underlying space into regions, such that the points within one region have the same farthest site (with respect to a given convex polyhedral distance function, in our case). We will show that the complexity of this diagram in $\mathbb{R}^d$ is $\Theta(n^{\lceil \frac{d}{2} \rceil})$ in the worst-case, and that it can be computed in optimal time $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$, mainly by using higher-dimensional convex hull algorithms. This result holds under rather general conditions: Sites can be arbitrary convex polyhedra (which may be unbounded or overlapping, having a total of n faces of various dimensions), the distance polytope $\mathcal{P}$ may be unbounded (though constant-sized), and the resulting distance function can be additively and/or multiplicatively weighted for each site. Moreover, each site may get allotted a particular distance polytope, in order to generate an anisotropic scenario where sites can influence their surrounding in an individual way.

Farthest-site Voronoi diagrams are useful for performing farthest neighbor queries among the sites, for computing the smallest ball that contacts all sites, and for finding the largest gap to be bridged between any two sites – to name a few or their applications. Unfortunately, Euclidean farthest-site Voronoi diagrams have their peculiarities (unless all sites are points, in which case their combinatorial and computational behavior is much like their closest-point counterparts [20]). Their regions may disconnect into a large number of nonconvex parts, and the close relationship between nonempty regions and the convex hull of the sites is lost; see [3] for line segment sites in $\mathbb{R}^2$, and [19] for a generalization to arbitrary L_p-metrics. The only result for non-point sites in higher dimensions we are aware of is [5], who derive structural and

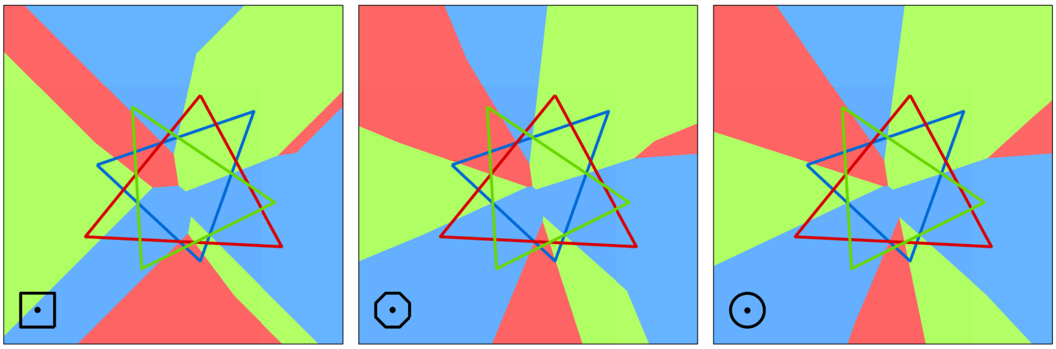

Figure 1 Two approximations (a) and (b) of a Euclidean farthest-site Voronoi diagram (c). The sites are three overlapping triangles. Their boundaries are visualized in individual colors, and their farthest regions are painted accordingly. The distance polygons used – a square in (a) and a regular 8-gon in (b) – are shown in the bottom-left corner.

combinatorial properties for the farthest-site diagram of lines and line segments in $\mathbb{R}^d$. They characterize its unbounded cells, of which there are up to $\Theta(n^{d-1})$ many in the worst-case, and describe an algorithm to compute these in near-optimal time.

With our results in the present note, a large class of Euclidean farthest-site Voronoi diagrams for convex sites in $\mathbb{R}^d$, even in their weighted and/or anisotropic variants, can be approximated in a piecewise-linear manner, and are computable by a simple and uniform approach: In $\mathbb{R}^3$ for example, being probably the most interesting case, the Euclidean ball can be β-approximated by a convex polytope $\mathcal{P}$ with $O(1/\beta)$ vertices [16], such that the convex distance induced by $\mathcal{P}$ is at most $1+\beta$ times the Euclidean distance. As a particularly useful result, a simple method for computing a piecewise-linear approximation of size $O(n^2)$ of the Euclidean farthest-site Voronoi diagram for lines and/or line segments in $\mathbb{R}^3$ becomes available. Even the planar instance is interesting: The fastest known algorithm for the Euclidean farthest-site Voronoi diagram for polygonal sites in $\mathbb{R}^2$ runs in time $O(n \log^3 n)$ [8], whereas our approximation can be computed in time $O(n \log n)$ if the polygonal sites are convex. Figure 1 illustrates the similarity between these diagrams.

2 Convex polyhedral distance

We define a *polyhedron* in $\mathbb{R}^d$ as the nonempty and finite (but possibly unbounded) intersection of closed halfspaces of $\mathbb{R}^d$. Note that a polyhedron does not need to be full-dimensional: For example lines, line segments, and single points are included as lower-dimensional instances.

Any d-dimensional polyhedron $\mathcal{P}$ which contains the origin in its interior can be used to define a so-called *convex polyhedral distance*, from a point $x \in \mathbb{R}^d$ to a point $q \in \mathbb{R}^d$:

$$\delta_{\mathcal{P}}(x, q) = \inf_{t \geq 0} \{ t \mid q \in x + t \cdot \mathcal{P} \} .$$

In other words, $\delta_{\mathcal{P}}(x, q)$ describes the amount $t \geq 0$ by which $\mathcal{P}$, when being placed at x, has to be scaled so as to cover q; see Figure 2. Note that $\delta_{\mathcal{P}}$ is a directed distance. We shall call $\mathcal{P}$ the *distance polytope* that induces the distance function $\delta_{\mathcal{P}}$.

Let $\mathcal{P}^R = \{ -p \mid p \in \mathcal{P} \}$ denote the reflection of the distance polytope about the origin.

▶ **Observation 1.** *We have* $\delta_{\mathcal{P}}(x,q) = \delta_{\mathcal{P}^R}(q,x)$.

Proof. Suppose that $q \in x + t\cdot\mathcal{P}$. Then there is a point $p \in \mathcal{P}$ with $q = x + t\cdot p$, that is, $x = q - t\cdot p$. Thus we have $x \in q - t\cdot\mathcal{P}$, which by the identity $-t\cdot\mathcal{P} = t\cdot\mathcal{P}^R$ means $x \in q + t\cdot\mathcal{P}^R$. ◀

Consider a set S of point sites in $\mathbb{R}^d$, and identify $\mathbb{R}^d$ with the hyperplane $x_{d+1} = 0$ in $(d+1)$-dimensional space $\mathbb{R}^{d+1}$. Observation 1 suggests to associate the distance polytope $\mathcal{P}$ with a *distance cone* $C_{\mathcal{P}}$ in $\mathbb{R}^{d+1}$, such that $C_{\mathcal{P}}$ reflects with its height the polyhedral distance induced by $\mathcal{P}$.

$$C_{\mathcal{P}} = \bigcup_{t\geq 0} \begin{pmatrix} t\cdot\mathcal{P}^R \\ t \end{pmatrix} \tag{1}$$

$C_{\mathcal{P}}$ is a polyhedral cone obtained from scaling the reflected polytope $\mathcal{P}^R$. Its apex is at the origin. Let $C_{\mathcal{P}}(q_i)$ be the translate of $C_{\mathcal{P}}$ with its apex at some point site $q_i \in S$. Then for any point $x \in \mathbb{R}^d$, the $(d+1)^{st}$ coordinate (called *height*) of the vertical projection of x to $C_{\mathcal{P}}(q_i)$ equals the distance $\delta_{\mathcal{P}}(x, q_i)$.

Let now, more generally, the set S consist of polyhedral sites s_i in $\mathbb{R}^d$. We construct for each site $s_i \in S$ a distance cone as follows. Take the Minkowski sum $s_i \oplus C_{\mathcal{P}}$. (The *Minkowski sum* of point sets A and B is defined as $A \oplus B = \{\, a+b \mid a \in A \wedge b \in B \,\}$.) Because the Minkowski sum of two convex polyhedra is again a convex polyhedron, the object $s_i \oplus C_{\mathcal{P}}$ is the intersection of halfspaces of $\mathbb{R}^{d+1}$. One of them is bounded from below by the hyperplane $x_{d+1} = 0$ (if the site s_i is full-dimensional). We ignore this halfspace, and intersecting the remaining ones we obtain an unbounded polyhedron in $\mathbb{R}^{d+1}$, which we call the *distance cone* of s_i, and denote with $C_{\mathcal{P}}(s_i)$.

$C_{\mathcal{P}}(s_i)$ exhibits the following useful properties.

- For the special case of s_i being a point site q_i, the definition of $C_{\mathcal{P}}(s_i)$ is consistent with that of $C_{\mathcal{P}}(q_i)$ before.
- Let $d_{\mathcal{P}}(x, s_i)$ be the height of the vertical projection of a point $x \in \mathbb{R}^d$ to $C_{\mathcal{P}}(s_i)$. If x does not lie in the interior of s_i, then $d_{\mathcal{P}}(x, s_i)$ is non-negative and equals the polyhedral distance of x to s_i, which is commonly defined as

  $$\delta_{\mathcal{P}}(x, s_i) = \inf_{t\geq 0} \{\, t \mid x + (t\cdot\mathcal{P}) \cap s_i \neq \emptyset \,\}\,.$$

- If x lies in the interior of s_i then $d_{\mathcal{P}}(x, s_i)$ is negative, and measures how much x is inside the (full-dimensional) polyhedral site s_i by taking the minimum polyhedral distance to its facets; see Figure 1. This is because the part of $C_{\mathcal{P}}(s_i)$ that lies below the hyperplane $x_{d+1} = 0$ is determined solely by halfspaces which stem from s_i (and not from $C_{\mathcal{P}}$ in Formula (1)). That is, $d_{\mathcal{P}}$ is related to a generalized medial axis of s_i in this case.

3 Farthest-site Voronoi diagram

The so-called *farthest-site Voronoi diagram* of a set S of sites in $\mathbb{R}^d$, for short $\mathrm{FVD}(S)$, is a partition of $\mathbb{R}^d$ into regions such that all points within a fixed region have the same farthest site. As before, let the sites s_i in S be polyhedra. These may be of any dimension k, for $0 \leq k \leq d$, and are not required to be disjoint or bounded.

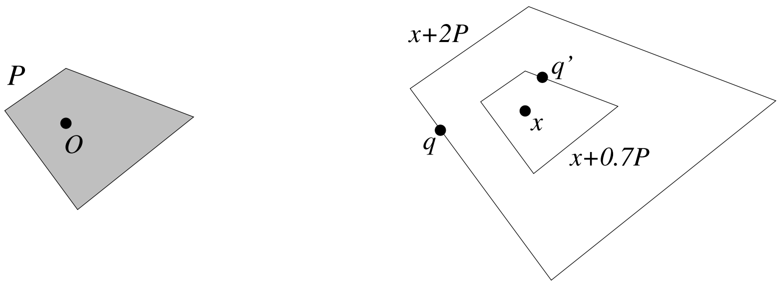

Figure 2 Polyhedral distance induced by $\mathcal{P}$: $d_{\mathcal{P}}(x,q) = 2$ and $d_{\mathcal{P}}(x,q') = 0.7$.

We are interested in the diagram FVD(S) induced by the convex polyhedral distance function $d_{\mathcal{P}}$ in Section 2, for a given distance polytope $\mathcal{P}$. Being a farthest-site diagram, FVD(S) corresponds to the pointwise maximum of the functions $d_{\mathcal{P}}(x, s_i)$, for $s_i \in S$, on R^d. FVD(S) thus corresponds to the *upper envelope* of the boundaries of the distance cones $C_{\mathcal{P}}(s_i)$ that define these functions, which, in turn, is given by the common intersection of these cones. Let us formulate this result in the following way.

▶ **Theorem 2.** *Let I be the (unbounded) convex polyhedron in $\mathbb{R}^{d+1}$ that results from intersecting the distance cones $C_{\mathcal{P}}(s_i)$, for all sites $s_i \in S$. Then* FVD(S) *is the vertical projection of I onto the hyperplane $x_{d+1} = 0$ of $\mathbb{R}^{d+1}$.*

One of the consequences of Theorem 2 is that FVD(S) is a piecewise linear diagram. Each region of FVD(S) is pre-partitioned into convex polyhedra (the projected facets of I), and these regions define a partition of $\mathbb{R}^d$. Let us point out that, in earlier papers on Voronoi diagrams for polyhedral distance functions (e.g. in [16]), the distance of a point x to a site was set to zero in case x falls in the interior of that site. As a consequence, when the sites are not chosen to be pairwise disjoint, the part of $\mathbb{R}^d$ covered by their union does not get partitioned by the diagram. Our more general definition of polyhedral distance, via distance cones, remedies this shortcoming.

The combinatorial complexity of FVD(S) is given by that of the projection polyhedron I in $\mathbb{R}^{d+1}$. I is the intersection of distance cones, and each distance cone $C_{\mathcal{P}}(s_i)$, in turn, is the intersection of halfspaces of $\mathbb{R}^{d+1}$. It is clear from Section 2 that the number of such halfspaces per cone is bounded by the number of facets of the Minkowski sum $s_i \oplus \mathcal{P}^R$, for the reflected distance polytope $\mathcal{P}^R$. A single face of s_i, combined with a single face of $\mathcal{P}^R$, can yield at most one facet of $s_i \oplus \mathcal{P}^R$; see e.g. [10]. Therefore, if we assume that $\mathcal{P}$ (and with it $\mathcal{P}^R$) is of constant size, and that s_i has a total of n_i faces of different dimensions, then $C_{\mathcal{P}}(s_i)$ is defined by $O(n_i)$ halfspaces of $\mathbb{R}^{d+1}$. In conclusion, when putting $n = \sum_{s_i \in S} n_i$, the polyhedron I is the intersection of $O(n)$ halfspaces of $\mathbb{R}^{d+1}$, and its complexity is bounded from above by $O(n^{\lceil \frac{d}{2} \rceil})$ (provided $d = O(1)$), by the well-known upper bound theorem. We will show in Section 4 that this complexity can be asymptotically attained in the worst case. Observe that I has $O(n)$ facets, and that FVD(S) thus has this very number of full-dimensional cells.

Concerning computational aspects, the halfspaces defining a particular cone $C_{\mathcal{P}}(s_i)$ can be singled out by (basically) computing the Minkowski sum $s_i \oplus \mathcal{P}^R$. This can be done [10], for instance, by pairwise adding up the $O(n_i)$ vertices of s_i and the $O(1)$ vertices of $\mathcal{P}^R$, and computing the convex hull of the resulting $O(n_i)$ points in $\mathbb{R}^d$, spending a total of $O(n^{\lfloor \frac{d}{2} \rfloor} + n \log n)$ time for all sites $s_i \in S$, when the optimal convex hull algorithm in [7] is applied. The construction of the projection polyhedron I is more time-consuming and takes $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$ time; we use the convex hull algorithm in [7] again, but now for intersecting $O(n)$ halfspaces in $\mathbb{R}^{d+1}$.

We may summarize as follows:

▶ **Theorem 3.** *Let S be a set of arbitrary polyhedral sites in $\mathbb{R}^d$, with a total combinatorial complexity of n. The farthest-site Voronoi diagram* $\mathrm{FVD}(S)$ *of S under the convex distance function induced by a polytope of constant size is of complexity $O(n^{\lceil \frac{d}{2} \rceil})$, and it can be computed in $O(n^{\lceil \frac{d}{2} \rceil} + n \log n)$ time. The number of d-dimensional cells of* $\mathrm{FVD}(S)$ *is bounded by $O(n)$.*

The dependence on d of the bounds stated above is the same as for convex hulls of finite point sets.

4 More properties of FVD

The maximal size of farthest-site diagrams may be much smaller than that of their closest-site counterparts; several examples can be found in [4]. The question arises whether the upper bound given in Theorem 3 is asymptotically tight.

For special sets of polyhedral sites in $\mathbb{R}^d$, the diagram $\mathrm{FVD}(S)$ is indeed small, namely, when the sites in S have only a constant number of orientations. Then the halfspaces defining the projection polyhedron I in $\mathbb{R}^{d+1}$ will have only a constant number of orientations as well, and all but $O(1)$ of them will be redundant because their bounding hyperplanes are parallel. Consequently, the polyhedron I and its projection $\mathrm{FVD}(S)$ will be of constant size, and can be found in $O(n)$ time.

Observe that the case of S being a set of n *point sites* in $\mathbb{R}^d$ is covered above, because each point site can be described by the intersection of $d+1$ halfspaces of $\mathbb{R}^d$, having the same fixed orientations. Not included, however, is the case of n line segment sites in $\mathbb{R}^d$, because the $d+2$ halfspaces describing a line segment will be of different orientation for different sites, in general. In fact, sites of very simple shape can induce large diagrams, as is shown below.

▶ **Lemma 4.** *There exists a set S of n sites in $\mathbb{R}^d$ of constant description such that the diagram* $\mathrm{FVD}(S)$ *has a complexity of $\Omega(n^{\lceil \frac{d}{2} \rceil})$.*

Proof. There exist two hyperplanes h_1, h_2 in $\mathbb{R}^{d+1}$, and two point sets $Y_1 \subset h_1$ and $Y_2 \subset h_2$ each of size $\frac{n}{2}$, such that the lower convex hull of $Y_1 \cup Y_2$ has $\Theta(n^{\lfloor \frac{d+1}{2} \rfloor}) = \Theta(n^{\lceil \frac{d}{2} \rceil})$ complexity; this follows from Corollary 12 in [14]. W.l.o.g., we may assume (by applying an affine transformation) that

$$\begin{aligned} &h_1 : \sum_{i=1}^{d} x_i = 1, \quad h_2 : \sum_{i=1}^{d-1} x_i - x_d = 1, \quad \text{and} \\ &Y_1 \subset \mathbb{R} \times \left(0, \frac{1}{d}\right]^d, \quad Y_2 \subset \mathbb{R} \times \left(0, \frac{1}{d}\right]^{d-1} \times \left[-\frac{1}{d}, 0\right), \end{aligned} \tag{2}$$

which fixes the hyperplanes and guarantees that most coordinates in $Y_1 \cup Y_2$ are small. We now choose a distance polytope $\mathcal{P}$ and a set $S = S_1 \cup S_2$ of sites such that the projection polyhedron I for $\mathrm{FVD}(S)$ is dual to the lower convex hull of $Y_1 \cup Y_2$.

Let the hypercube $[-1,1]^d$ serve as $\mathcal{P}$. The set $S_1 \cup S_2$ will consist of n halfspace sites in $\mathbb{R}^d$. For S_1, each of its halfspaces s is constructed from a point $y = (y_1, \ldots, y_{d+1}) \in Y_1$. In particular, we describe s by the inequality $\sum_{i=1}^d a_i x_i \leq b$, where

$$a_1 = 1, \quad a_i = \frac{y_i}{1 - \sum_{j=2}^d y_j} \quad \text{for} \quad i = 2, \ldots d, \quad \text{and} \quad b = \frac{y_{d+1}}{1 - \sum_{j=2}^d y_j}. \tag{3}$$

Note that all a_i and b are positive because of our assumption (2). Moreover, we have $\mathcal{P}^R = \mathcal{P}$. Therefore, the Minkowski sum $s \oplus \mathcal{P}^R$ is just a translate of s by the vector $(1, 1, \ldots, 1)^T$ in $\mathbb{R}^d$, which implies that the distance cone $C_{\mathcal{P}}(s)$ is a single halfspace in $\mathbb{R}^{d+1}$, bounded from below by the hyperplane

$$x_{d+1} = \frac{1}{A} \cdot \left(\sum_{i=1}^d a_i x_i - b \right), \quad \text{where} \quad A = \sum_{j=1}^d a_j. \tag{4}$$

By a well-known duality transform, the hyperplane in (4) is dual to the point

$$q = (q_1, \ldots, q_{d+1}) = \frac{1}{A} \cdot (a_1, \ldots, a_d, b)$$

in $\mathbb{R}^{d+1}$. Substituting the values in (3) and simple calculations give

$$q_i = y_i \quad \text{for} \quad i = 2, \ldots, d+1, \quad \text{and} \quad q_1 = 1 - \sum_{j=2}^d y_j. \tag{5}$$

But this implies $q = y$, because both points lie in the hyperplane h_1 in (2): We have $y \in h_1$ by assumption, and $q \in h_1$ by (5).

In a similar manner, we can construct suitable halfspace sites s for S_2 from the points y in Y_2. (We omit these details here.) In conclusion, the projection polyhedron I, being the intersection of all the sites' distance cones, is the intersection of the upper halfspaces $C_{\mathcal{P}}(s)$ for all $s \in S_1 \cup S_2$, and thus I is dual to the lower convex hull of $Y_1 \cup Y_2$. ◀

By Lemma 4, the runtime in Theorem 3 is asymptotically optimal in the worst case for $d \geq 3$. For $d = 2$, a reduction from sorting proves optimality.

Being the projection of I, the diagram FVD(S) is a polyhedral cell complex in $\mathbb{R}^d$ which is face-to-face. Its *cells* (polyhedra of dimension d) are nonconvex in general, as are its *facets* (polyhedra of dimension $d-1$). Since the distance polytope $\mathcal{P}$ is (more or less) an approximation of the Euclidean ball, quite a few properties of the Euclidean farthest-site diagram of S carry over to FVD(S); see e.g. [3, 5]. For example, the *region* of a site $s_i \in S$ in FVD(S) (the set of all points in $\mathbb{R}^d$ being farthest from s_i) is disconnected in general, and it may consist of various cells of FVD(S). Moreover, the following properties of the cells are preserved.

▶ **Lemma 5.** *All cells of* FVD(S) *are unbounded, and cells cannot contain voids of any dimension.*

Proof. Let C be some cell of FVD(S), and assume that C is part of the region of the site $s \in S$. The assertion of the lemma can be easily derived from the following fact: Let x be an arbitrary point in C, and consider the point y on the boundary of s that realizes the polyhedral distance $d_{\mathcal{P}}(x, s)$. Then the infinite ray r that starts at x and is directed away from y is totally contained in C.

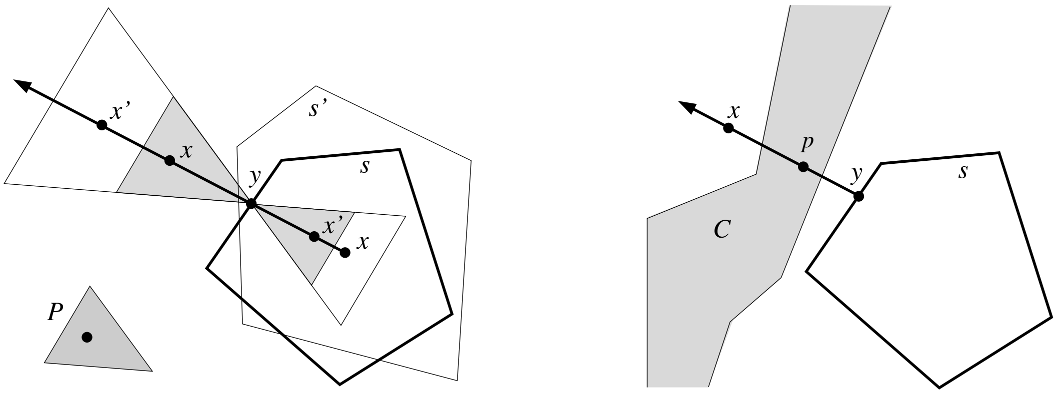

Figure 3 Illustrations of the proofs of Lemma 5 (left) and Lemma 6 (right).

To prove this fact, refer to Figure 3 (left). Assume first that $x \notin s$. Then $t = d_{\mathcal{P}}(x, s) \geq 0$, and the homothet $H = x + t \cdot \mathcal{P}$ touches s at y. Since s is the site in S farthest from x, H intersects all the other sites. Let now x' be any point on r such that x lies between x' and y. Put $t' = d_{\mathcal{P}}(x', s)$. Then $H' = x' + t' \cdot \mathcal{P}$ touches s at y too, and H is covered by H', which implies that H' intersects all other sites as well. This implies that x' lies in the region of s.

If $x \in s$, on the other hand, then $t = d_{\mathcal{P}}(x, s) < 0$, and we have $t \cdot \mathcal{P} = u \cdot \mathcal{P}^R$ with $u = -t > 0$, for the reflected polytope $\mathcal{P}^R$. The homothet $H = x + u \cdot \mathcal{P}^R$ touches s at y, and since s is farthest from x, H is contained in all other sites now. For any point x' on r between x and y, and $u' = -d_{\mathcal{P}}(x', s)$, $H' = x' + u' \cdot \mathcal{P}^R$ touches s at y again, but is contained in H now and therefore also in all other sites. So x' has to lie in the region of s.

In summary, we conclude that the entire ray r lies in the cell C of the region of s. ◀

Let us define the $(d-1)$-*skeleton* of FVD(S) as the union of all the facets of FVD(S). This skeleton can be disconnected, as a simple construction with only two sites shows; see Figure 4(a): Let site s_1 be some polyhedron which approximates a line segment, and take as site s_2 any polyhedron which contains the segment's midpoint but none of its endpoints. Then the region of s_1 disconnects the $(d-1)$-skeleton of FVD($\{s_1, s_2\}$). On the other hand, by the same argument as in [5], the following holds:

▶ **Lemma 6.** *The $(d-1)$-skeleton of* FVD(S) *is connected, provided that the sites in S are pairwise disjoint.*

Proof. Assume that this skeleton is not connected; see Figure 3 (right). Then there exists some cell C of FVD(S) that splits the skeleton into at least two parts. Let s be the farthest site corresponding to C. The site s does not touch the boundary of C, because of our assumption on the disjointness of the sites. Thus there exists some point $x \notin C$ which is separated from s by C. Let y be the point on s that realizes the polyhedral distance from x to s. By construction, the line segment $\overline{xy}$ intersects C, and we choose a point p in this intersection. Now, by the reasoning in the proof of Lemma 5, the infinite ray emanating from p in direction x is entirely contained in C. But this implies $x \in C$, which is a contradiction. ◀

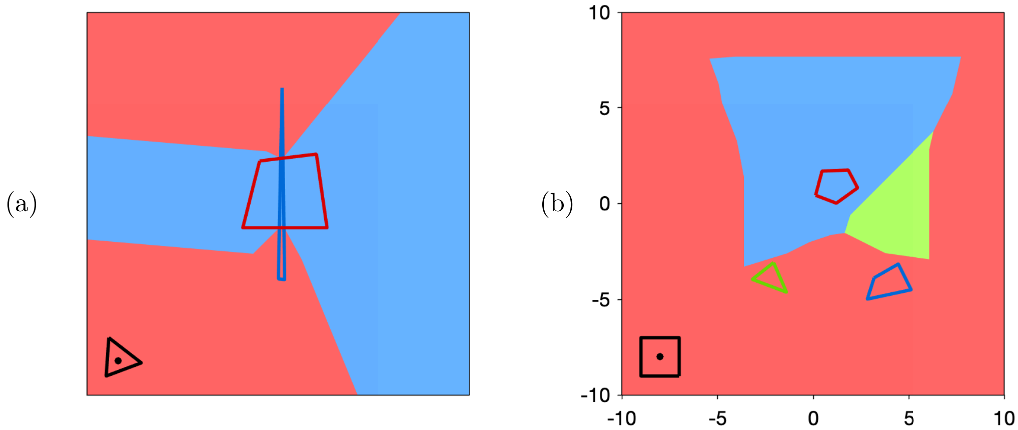

Figure 4 (a) The $(d-1)$-skeleton can be disconnected for non-disjoint sites. (b) A weighted farthest Voronoi diagram of three sites: The blue quadrangle has an additive weight of -1, and the red pentagon has a multiplicative weight of $\frac{1}{2}$.

5 Variants

In certain applications, a model of Voronoi diagram is required where the sites are capable of influencing their surrounding in an individual way; see [4, 18] for comprehensive treatments of this topic. One way to achieve this goal is to assign so-called *weights* to the sites, which affect the underlying distance function in an additive and/or multiplicative way.

Let each site $s_i \in S$ have assigned two real numbers $a(s_i)$ and $m(s_i) > 0$, and consider the weighted polyhedral distance:

$$\frac{d_{\mathcal{P}}(x, s_i)}{m(s_i)} - a(s_i)$$

In contrast to the nearest version, the sites' regions in the farthest Voronoi diagram shrink with increasing weights. Interestingly, and unlike the situation for the Euclidean farthest-site diagram, the $\mathrm{FVD}(S)$ induced by this distance is still a piecewise-linear cell complex. This becomes evident when the respective distance cones are considered: Additive weighting results in a vertical shift of these cones by an amount of $a(s_i)$, and multiplicative weighting enlarges by a factor of $m(s_i)$ the value $\tan \alpha_j$ of the dihedral angles α_j of aperture of a cone's facets. In particular, each distance cone still is the intersection of $O(n_i)$ halfspaces of $\mathbb{R}^{d+1}$ when site s_i is of complexity n_i.

Multiplicative weighting leads to the occurrence of bounded regions in $\mathrm{FVD}(S)$, as simple examples show (Figure 4 (b)). However, purely additive weighting preserves the properties listed in Lemma 5. In particular, all cells are still unbounded: All facets of the projection polyhedron I for the unweighted $\mathrm{FVD}(S)$ are unbounded, and this fact cannot be altered by vertically shifting any of its defining halfspaces.

We may push things even further, and create an *anisotropic* scenario by allotting an individual distance polytope $\mathcal{P}_i$ to each site s_i. In this way, each site is able to "interpret" its surrounding space in its own way – a concept useful in many situations [4]. In fact, the multiplicative weighting scheme is just the special case where $\mathcal{P}_i = m(s_i) \cdot \mathcal{P}$.

In all the extensions above, the properties of the distance cones needed for Theorem 3 to hold are preserved. We obtain the following general result:

▶ **Theorem 7.** *Theorem 3 remains valid for all the weighted and anisotropic variants of* $\mathrm{FVD}(S)$ *described above.*

Note finally that all these extensions can be combined, and lead to a very general class of easy-to-compute piecewise-linear farthest-site Voronoi diagrams in $\mathbb{R}^d$, where the impact of each site can be tuned by its shape, its weights, and its distance polytope including the choice of the polytope's center.

References

1 Boris Aronov. A lower bound on Voronoi diagram complexity. *Information Processing Letters*, 83(4):183–185, 2002.

2 Franz Aurenhammer. Power diagrams: Properties, algorithms and applications. *SIAM Journal on Computing*, 16(1):78–96, 1987.

3 Franz Aurenhammer, Scot Drysdale, and Hannes Krasser. Farthest line segment Voronoi diagrams. *Information Processing Letters*, 100(6):220–225, 2006.

4 Franz Aurenhammer, Rolf Klein, and Der-Tsai Lee. *Voronoi diagrams and Delaunay Triangulations*. World Scientific Publishing Company, 2013.

5 Gill Barequet, Evanthia Papadopoulou, and Martin Suderland. Unbounded regions of high-order Voronoi diagrams of lines and segments in higher dimensions. In *30th International Symposium on Algorithms and Computation*, volume 149 of *LIPIcs*, pages 62:1–62:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

6 Jean-Daniel Boissonnat, Micha Sharir, Boaz Tagansky, and Mariette Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete & Computational Geometry*, 19(4):485–519, 1998.

7 Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10:377–409, 1993.

8 Otfried Cheong, Hazel Everett, Marc Glisse, Joachim Gudmundsson, Samuel Hornus, Sylvain Lazard, Mira Lee, and Hyeon-Suk Na. Farthest-polygon Voronoi diagrams. *Computational Geometry: Theory and Applications*, 44(4):234–247, 2011.

9 L. Paul Chew, Klara Kedem, Micha Sharir, Boaz Tagansky, and Emo Welzl. Voronoi diagrams of lines in 3-space under polyhedral convex distance functions. *Journal of Algorithms*, 29(2):238–255, 1998.

10 Sandip Das and Swami Sarvottamananda. Computing the Minkowski sum of convex polytopes in R^{d}. *CoRR*, abs/1811.05812, 2018. `arXiv:1811.05812`.

11 Herbert Edelsbrunner and Raimund Seidel. Voronoi diagrams and arrangements. *Discrete & Computational Geometry*, 1:25–44, 1986.

12 Hazel Everett, Daniel Lazard, Sylvain Lazard, and Mohab Safey El Din. The Voronoi diagram of three lines. *Discrete & Computational Geometry*, 42(1):94–130, 2009.

13 Christian Icking and Lihong Ma. A tight bound for the complexity of Voroni diagrams under polyhedral convex distance functions in 3D. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing*, pages 316–321, 2001.

14 Menelaos I. Karavelas, Raimund Seidel, and Eleni Tzanaki. Convex hulls of spheres and convex hulls of disjoint convex polytopes. *Computational Geometry: Theory and Applications*, 46(6):615–630, 2013.

15 Victor Klee. On the complexity of d-dimensional Voronoi diagrams. *Archiv der Mathematik*, 34(1):75–80, 1980.

16 Vladlen Koltun and Micha Sharir. Polyhedral Voronoi diagrams of polyhedra in three dimensions. *Discrete & Computational Geometry*, 31(1):83–124, 2004.

17 Joseph S. B. Mitchell and Joseph O'Rourke. Computational Geometry Column 42. *Intern. Journal of Computational Geometry & Applications*, 11(5):573–582, 2001.

18 Atsuyuki Okabe, Barry Boots, Kokichi Sugihara, and Sung Nok Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams (2nd Edition)*, volume 501. John Wiley & Sons, 2009.

19 Evanthia Papadopoulou and Sandeep Kumar Dey. On the farthest line-segment Voronoi diagram. *Intern. Journal of Computational Geometry & Applications*, 23(6):443–460, 2013.

20 Raimund Seidel. On the number of faces in higher-dimensional Voronoi diagrams. In *Proceedings of the Third Annual Symposium on Computational Geometry*, pages 181–185, 1987.

21 Micha Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete & Computational Geometry*, 12:327–345, 1994.