

Voronoi Diagrams and Arrangements

Herbert Edelsbrunner^{1*} and Raimund Seidel²

¹Institute f. Informationsverarbeitung, Technical University of Graz, Schießstattgasse 4A, A-8010 Graz, Austria

²Computer Science Department, Cornell University, Ithaca, NY 14853

Abstract. We propose a uniform and general framework for defining and dealing with Voronoi diagrams. In this framework a Voronoi diagram is a partition of a domain D induced by a finite number of real valued functions on D . Valuable insight can be gained when one considers how these real valued functions partition $D \times \mathbf{R}$. With this view it turns out that the standard Euclidean Voronoi diagram of point sets in \mathbf{R}^d along with its order- k generalizations are intimately related to certain arrangements of hyperplanes. This fact can be used to obtain new Voronoi diagram algorithms. We also discuss how the formalism of arrangements can be used to solve certain intersection and union problems.

1. Introduction

Figure 1.1 depicts a diagram of a type known as Dirichlet tessellation, Thiessen polygons, or as we call it, Voronoi diagram. The formation rule for such a diagram is simple. The location of a finite number of "sites" is known. For each "site" s one wants to form the region of all points for which s is the nearest among the finite set of "sites." If "nearest" is understood with respect to the Euclidean distance measure, then for each "site" its associated region is polygonal.

In the context of computational geometry Voronoi diagrams were first introduced in a paper by Hoey and Shamos [31]. The usefulness of the Voronoi diagram (referred to from now on as VoD) for solving a large number of problems, the fact that it can be constructed efficiently, and maybe also its

*Present address: Department of Computer Science, University of Illinois, 1304 West Springfield Ave., Urbana, IL 61801.




Fig. 1.1. A Voronoi Diagram.

aesthetically pleasing appearance subsequently kindled the interest of many researchers. They tried to apply the VoD to other problems, and if this was not possible directly, attempted to adapt and generalize the notion of a VoD appropriately. Among the generalizations are higher order VoDs [21], [31], VoDs of line segments and/or circular arcs [12], [20], [23], [32], [34], VoDs of point sets in \mathbb{R}^d , $d > 2$ [5], [33], weighted Voronoi diagrams [3], power diagrams, or VoDs with respect to the Laguerre geometry [1], [19], VoDs with respect to the L^p metric, $1 \leq p \leq \infty$ [22], [24], and VoDs with respect to other “funny” metrics [2], [9].

The types of VoDs just mentioned are quite different and there is little doubt that even more generalizations will be investigated in the future. However, instead of pursuing ever more diversification one can also attempt unification and ask: What do all these different types of VoDs have in common: What constitutes a VoD in its most general form? What are the underlying mathematical concepts? Does one, for instance, need the notion of a metric in order to define VoDs?

In this paper we try to answer some of these questions. We offer a very general definition of VoDs which shows that they can be defined naturally on any domain D via a finite set of real valued functions on D . Thus no notion of metric is needed. We show that higher order and higher degree VoDs can be obtained in a similar manner and demonstrate how all these VoDs are related to the partition of the product space $D \times \mathbb{R}$ induced by the finite set of real valued functions on D .

This view already proves useful when applied to the standard Euclidean VoD of point sets in \mathbb{R}^d . It turns out that this kind of VoD, along with its higher order and higher degree versions, can be defined using a finite set of affine functions from \mathbb{R}^d to \mathbb{R} . Consequently these diagrams are intimately related to arrangements of hyperplanes in \mathbb{R}^{d+1} , an insight which leads to new algorithms for Euclidean VoDs.

Section 2 of this paper deals with arrangements, i.e., partitions of $D \times \mathbb{R}$ induced by a finite number of real valued functions on D . In Section 3 we show how VoDs can be defined on any domain D via finite sets of real valued functions. We demonstrate how VoDs are related to arrangements and we discuss a number of examples in some detail. Section 4 deals with higher order and higher degree VoDs and their relationship to arrangements which turns out to be particularly attractive. Section 5 contains a short outline of how the formalism of arrangements can be used to solve certain intersection and union problems. In the last section we discuss possible directions for further research.

2. Arrangements

In this section we consider the interactions among a finite number of real valued functions on an arbitrary domain D . The structure imposed on $D \times \mathbf{R}$ by such a finite set of functions promises to be an interesting object of study, provided the domain D and the functions as well as their interactions are restricted appropriately. For instance, when $D = \mathbf{R}^d$ and all the functions are affine, the structure of $D \times \mathbf{R}$ has been studied to a fair extent as so-called hyperplane arrangement [17], [18], [35]. Similarly, in a slight generalization, when $D = \mathbf{R}^d$ and all functions are continuous and satisfy certain finite intersection axioms, the structure imposed on $D \times \mathbf{R}$ has been studied as so-called arrangement of pseudo-hyperplanes, and also in the context of oriented matroids [26].

In this section we do not intend to study the effect of any other function or domain restriction. We rather want to give some general definitions and results in order to provide a convenient framework to argue about Voronoi diagrams, the main topic of this paper.

Let D be some domain. For a real valued function f on D we call the subset of $D \times \mathbf{R}$

$$f^+ = \{(x, r) | r < f(x)\}$$

the *lower hemispace* of f , the set

$$f^- = \{(x, r) | r > f(x)\}$$

the *upper hemispace* of f , and, with a slight abuse of terminology,

$$f^0 = \{(x, r) | r = f(x)\}$$

the *surface* of f . Note the f^+ , f^- , and f^0 properly partition $D \times \mathbf{R}$.

Throughout this paper let E denote a finite index set with n elements. For each $e \in E$ let f_e be a real valued function on D and let f_E denote the indexed collection of these n functions. For each point $y \in D \times \mathbf{R}$, let $\pi_f(y)$ denote the partition of E induced by the functions in f_E applied to point y , i.e.,

$$\pi_f(y) = \langle E^-, E^0, E^+ \rangle,$$

where

$$E^- = \{e \in E | y \in f_e^-\},$$

$$E^0 = \{e \in E | y \in f_e^0\},$$

$$E^+ = \{e \in E | y \in f_e^+\}.$$

In a natural way π_f defines an equivalence relation on $D \times \mathbf{R}$, making two points y and z equivalent iff $\pi_f(y) = \pi_f(z)$. We call the partition of $D \times \mathbf{R}$ induced by π_f an f_E -arrangement. We call each equivalence class of an f_E -arrangement a *cell*.




Fig. 2.1. The two-level of an arrangement.

Each cell can be uniquely named by a partition $\langle E^-, E^0, E^+ \rangle$ of E :

$$C_{\langle E^-, E^0, E^+ \rangle} = \{y \in D \times \mathbf{R} \mid \pi_f(y) = \langle E^-, E^0, E^+ \rangle\}.$$

Equivalently, a cell can be represented as an intersection:

$$C_{\langle E^-, E^0, E^+ \rangle} = \bigcap_{e \in E^-} f_e^- \cap \bigcap_{e \in E^0} f_e^0 \cap \bigcap_{e \in E^+} f_e^+.$$

As there are only a finite number of partitions of the index set E , the number of cells in an f_E -arrangement is finite and hence f_E -arrangements can be studied as combinatorial objects.

We call a cell $C_{\langle E^-, E^0, E^+ \rangle}$ a *full cell* iff $E^0 = \emptyset$. For an integer k , $0 \leq k \leq n$, we call a full cell $C_{\langle E^-, \emptyset, E^+ \rangle}$ a *k-belt cell* iff $|E^-| = k$. We call a cell $C_{\langle E^-, E^0, E^+ \rangle}$ a *k-level cell* iff $|E^-| < k$ and $|E^-| + |E^0| \geq k$. Note that a cell $C_{\langle E^-, E^0, E^+ \rangle}$ is in $|E^0|$ different levels. Figure 2.1 shows the two-level cells of an arrangement of three real valued functions on \mathbf{R} . Figure 2.2 depicts the two-belt cells of the same arrangement. Note that cells are not necessarily “connected.”

If f_E and g_E are two indexed collections of real valued functions on a common domain D , then we call f_E and g_E *order-equivalent* iff for all $x \in D$ and




Fig. 2.2. The two-belt of an arrangement.

for every pair $i, j \in E$

$$\text{sign}(f_i(x) - f_j(x)) = \text{sign}(g_i(x) - g_j(x)).$$

If two function collections f_E and g_E are order-equivalent, then the f_E -arrangement and g_E -arrangement of $D \times \mathbf{R}$ are isomorphic in the following sense:

Lemma 2.1. *Let f_E and g_E be two order-equivalent function collections on D .*

Let $\langle E^-, E^0, E^+ \rangle$ be a partition of E .

- (i) *There exists a nonempty cell $C_{\langle E^-, E^0, E^+ \rangle}$ in the f_E -arrangement iff there exists a nonempty cell $C'_{\langle E^-, E^0, E^+ \rangle}$ in the g_E -arrangement.*

Moreover,

- (ii) $\text{proj}(C_{\langle E^-, E^0, E^+ \rangle}) = \text{proj}(C'_{\langle E^-, E^0, E^+ \rangle})$, where for $C \subseteq D \times \mathbf{R}$ the expression $\text{proj}(C)$ denotes $\{x \in D \mid (x, r) \in C \text{ for some } r \in \mathbf{R}\}$.

Proof. To prove (i) and (ii) it suffices to show that for every $y = (x, r) \in D \times \mathbf{R}$ there exists an $r' \in \mathbf{R}$ such that $\pi_f((x, r)) = \pi_g((x, r'))$.

Let $\langle E^-, E^0, E^+ \rangle = \pi_f((x, r))$. Because of the order-equivalence between f_E and g_E we know that for $e^- \in E^-$, $e^+ \in E^+$, and $e, e' \in E^0$ the relations $g_{e^-}(x) < g_e(x) < g_{e^+}(x)$ and $g_e(x) = g_{e'}(x)$ hold.

Now let $l = \max\{g_e(x) \mid e \in E^- \cup E^0\}$, and $u = \min\{g_e(x) \mid e \in E^+ \cup E^0\}$. In case l does not exist choose $r' = u - 1$; if u does not exist choose $r' = l + 1$; and otherwise choose $r' = \frac{1}{2}(l + u)$. Such a choice of r' forces $\pi_f((x, r)) = \pi_g((x, r'))$. \square

In the remainder of this section we consider two important classes of arrangements. Let the domain D be \mathbf{R}^d for some $d \geq 1$. For $e \in E$ let f_e be an affine function given by

$$f_e(x) = \langle a_e, x \rangle + b_e,$$

where $a_e \in \mathbf{R}^d$, $b_e \in \mathbf{R}$, and $\langle \cdot, \cdot \rangle$ denotes the usual scalar product between vectors. The surface f_e^0 of such an affine function is usually called a nonvertical hyperplane in $\mathbf{R}^d \times \mathbf{R} = \mathbf{R}^{d+1}$. The upper and lower hemispaces f_e^- and f_e^+ are in this case usually referred to as upper and lower (open) halfspaces. The f_E -arrangement of \mathbf{R}^{d+1} is usually called an arrangement of hyperplanes, and is, as mentioned in the beginning of this section, a fairly well studied mathematical object. Every cell $C_{\langle E^-, E^0, E^+ \rangle}$ in an arrangement of hyperplanes is the intersection of a finite number of halfspaces and hyperplanes and hence a polyhedron. In the literature of hyperplane arrangements one usually considers the topological closure $\bar{C}_{\langle E^-, E^0, E^+ \rangle}$ of a cell $C_{\langle E^-, E^0, E^+ \rangle}$, with

$$\bar{C}_{\langle E^-, E^0, E^+ \rangle} = \bigcap_{e \in E^-} (f_e^- \cup f_e^0) \cap \bigcap_{e \in E^+} (f_e^+ \cup f_e^0) \cap \bigcap_{e \in E^0} f_e^0.$$

$\bar{C}_{\langle E^-, E^0, E^+ \rangle}$ is usually called a face of the arrangement. Faces are assigned a dimension, namely the dimension of the smallest flat containing the face. It is known that the number of faces is an arrangement of n hyperplanes in \mathbf{R}^{d+1} is

$O(n^{d+1})$. More specifically, let $c_i(n, d)$ denote the maximum possible number of faces of dimension i in an arrangement of n hyperplanes in \mathbb{R}^d .

Fact 2.1. [35].

$$c_i(n, d) = \sum_{d-i \leq j \leq d} \binom{j}{d-i} \binom{n}{j} \quad \text{for } 0 \leq i \leq d.$$

The algorithmic problem of identifying all cells of an arrangement of hyperplanes has been solved completely in [14].

Fact 2.2. [14]. A combinatorial description of all the faces of an arrangement of n hyperplanes in \mathbb{R}^{d+1} can be constructed in time and space $O(n^{d+1})$, which is asymptotically worst case optimal.

This description consists in essence of a directed graph representing the lattice formed by the faces under set inclusion. By increasing the space and time requirement by a constant factor this description can be extended such that for each face $\bar{C}_{\langle E^-, E^0, E^+ \rangle}$ its level or belt number can be determined in constant time and the index sets E^- , E^0 , or E^+ can be listed in time proportional to their respective sizes.

Not much is known about $l_k(n, d)$ and $b_k(n, d)$, the maximal possible number of k -level or k -belt faces in an arrangement of n hyperplanes in \mathbb{R}^d . Obviously $b_k(n, d) = b_{n-k}(n, d)$ and $l_k(n, d) = l_{n+1-k}(n, d)$ for all appropriate values of k . Furthermore, from the upper bound theorem for convex polytopes [25] and from the work in [13], [15], [16] the following is known:

Fact 2.3.

$$l_1(n, d+1) = b_1(n, d+1) = \Theta(n^{\lceil d/2 \rceil}).$$

For $k \leq n - k$

$$l_k(n, 2) = \Omega(n \log(k+1)),$$

$$l_k(n, 2) = O(n\sqrt{k}),$$

$$l_k(n, d+1) = \Omega\left(nk^{d-1} \log(k+1) + \frac{1}{k} n^{\min(\lceil d/2 \rceil, k)}\right),$$

$$l_k(n, d+1) = O(n^{d+1}).$$

The same asymptotic bounds hold for b_k as well.

As finding the one level faces is equivalent to constructing the polyhedron formed by the intersection of all lower halfspaces, [6], [28], [29], [30] imply

Fact 2.4. All one-level faces (or equivalently all n -level faces) of an arrangement of n hyperplanes in \mathbb{R}^{d+1} can be found in time

$$O(n \log n) \quad \text{for } d = 1, 2,$$

$$O(n^{\lceil (d+1)/2 \rceil}) \quad \text{for } d \geq 3.$$

This is worst case optimal for odd d and $d = 2$.

Much less is known about complexity bounds for constructing all k -level faces for arbitrary k .

Fact 2.5. [13]. All k -level faces of an arrangement of n planes in \mathbf{R}^3 can be found in time $O(\sqrt{k} \log nl_k(n,3))$.

This concludes our elaboration on the first class of examples of f_E -arrangements.

Our second class of examples deals with paraboloids. Let again the domain D of the functions be \mathbf{R}^d , $d \geq 1$. Let A be a $d \times d$ symmetric real matrix. For $e \in E$ let g_e be a quadratic function of the form

$$g_e(x) = (x - p_e)^T A (x - p_e) + t_e,$$

where $p_e \in \mathbf{R}^d$ and $t_e \in \mathbf{R}$. The surface g_e^0 is usually called a *paraboloid* and we thus call such g_E -arrangement in \mathbf{R}^{d+1} a *paraboloid arrangement*. To our knowledge paraboloid arrangements per se have not been studied in the literature. However, as the following important lemma shows, there is really no need to do so.

Lemma 2.2. For $e \in E$ let g_e be quadratic functions defined by

$$g_e(x) = (x - p_e)^T A (x - p_e) + t_e, \quad p_e \in \mathbf{R}^d \text{ and } t_e \in \mathbf{R}.$$

For $e \in E$ let f_e be affine functions defined by

$$f_e(x) = \langle x, a_e \rangle + b_e, \quad a_e = -2Ap_e, \quad b_e = p_e^T A p_e + t_e.$$

The collection of functions f_E and g_E are order-equivalent.

Proof. We have to show that for every $r \in \mathbf{R}^d$ and any pair $i, j \in E$

$$\text{sign}(f_i(x) - f_j(x)) = \text{sign}(g_i(x) - g_j(x)).$$

However, it can be checked easily that in our case even

$$\begin{aligned} f_i(x) - f_j(x) &= g_i(x) - g_j(x) = (p_i^T A p_i + t_i) \\ &\quad - (p_j^T A p_j + t_j) + 2(p_j - p_i)^T A x \end{aligned}$$

holds. \square

Lemmas 2.1 and 2.2 imply that the paraboloid arrangement and the hyperplane arrangement generated by g_E and f_E , respectively, are combinatorially indistinguishable with respect to their cell structure. As a matter of fact, the even stronger relation holds that every cell in the hyperplane arrangement is the homeomorphic image of the corresponding cell in the paraboloid arrangement

under the differentiable and invertible mapping F_A of \mathbf{R}^{d+1} onto itself defined by

$$F_A((x, z)) = (x, z - x^T A x), \quad \text{where } x \in \mathbf{R}^d \text{ and } z \in \mathbf{R}.$$

In other words, F_A "warps" \mathbf{R}^{d+1} in such a way that all paraboloids generated by the matrix A are flattened out into hyperplanes, but all intersection patterns are preserved. Thus all the results stated above about hyperplane arrangements apply to paraboloid arrangements as well.

3. Voronoi Diagrams

As in the previous section let E be a finite index set of n elements and let $f_E = \{f_e | e \in E\}$ be an indexed collection of real valued functions on some common domain D . Let R_f be a function from D to E defined by

$$R_f(x) = \left\{ e \in E | f_e(x) = \min_{i \in E} f_i(x) \right\}.$$

R_f induces in a canonical way an equivalence relation ρ_f on D , where for $x, y \in D$, $x \rho_f y \Leftrightarrow R_f(x) = R_f(y)$. We call the partition of D induced by this equivalence relation ρ_f the *Voronoi Diagram on D with respect to f_E* , for short $\text{VOD}(f_E)$. The equivalence classes of the partition are called *Voronoi cells* or *V-cells*. We denote each V-cell by V_T , where $T \subseteq E$ and

$$V_T = \{x \in D | R_f(x) = T\}.$$

A V-cell V_T with $|T|=1$ is sometimes also called a *Voronoi region* or *V-region*.

This general, functional definition of Voronoi diagrams might look somewhat startling and unorthodox to a reader only familiar with the usual definition of VoDs. Thus it seems appropriate to show that for the right choice of functions f_E , the Voronoi diagram of f_E obtained using our definition conforms with the standard Euclidean Voronoi diagram as presented, for instance, in [31].

Let D be \mathbf{R}^2 , the plane, and let $d(x, y)$ denote the Euclidean distance function. For each $e \in E$ let p_e be a point in the plane and define the real valued function f_e as

$$f_e(x) = d(p_e, x).$$

For some $e \in E$ consider the V-region $V_{\{e\}}$ in $\text{VOD}(f_E)$. $V_{\{e\}}$ contains all points $x \in \mathbf{R}^2$ with the property that

$$\left\{ e \in E | f_e(x) = \min_{i \in E} f_i(x) \right\} = \{e\},$$

or, in other words, $V_{\{e\}}$ contains all x with

$$f_e(x) < f_i(x) \quad \text{for all } i \in E, i \neq e,$$




Fig. 3.1.

i.e., all x with

$$d(p_e, x) < d(p_i, x) \quad \text{for all } i \in E, i \neq e.$$

Thus $V_{(e)}$ is exactly the interior of what is usually considered the Voronoi region of point p_e .

Similarly, if a V-cell $V_{(i,j)}$ exists in $\text{VOD}(f_E)$ for some pair $i, j \in E$, then this V-cell, by the same reasoning as above, contains all x with

$$d(p_i, x) = d(p_j, x) < d(p_e, x) \quad \text{for all } e \in E, i \neq e \neq j.$$

Thus $V_{(i,j)}$ is exactly the relative interior of the edge between the two V-regions $V_{(i)}$ and $V_{(j)}$.

Finally, V-cells of the form V_T with $|T| \geq 3$ turn out to be the vertices in the traditional VoD (see Fig. 3.1).

The reader may convince himself that with the right choice of functions other kinds of VoDs, such as Voronoi diagrams with respect to the L^p metric [22], VoDs of line segments [23], [20], [32], [34], and weighted VoDs [3], can be expressed using our functional formalism.

An almost trivial but very important observation is the fact that for a function collection f_E , the Voronoi diagram of f_E in D and the f_E -arrangement in $D \times \mathbf{R}$ are intimately related. We have the following:

Theorem 3.1. *Let f_E be a collection of real valued functions on a domain D , and let $\emptyset \neq T \subseteq E$, and $T' = E - T$. V_T is a V-cell in $\text{VOD}(f_E)$ iff $C_{\langle \emptyset, T, T' \rangle}$ is a cell in the f_E -arrangement. Moreover, $V_T = \text{proj}(C_{\langle \emptyset, T, T' \rangle})$. (Note that the cells $C_{\langle S, T, U \rangle}$ with $S = \emptyset \neq T$ are exactly the one-level cells.)*

Proof. By the definitions of this and the previous section, we have $x \in V_T$ if and only if for some $r \in \mathbf{R}$, $f_e(x) = r$ for $e \in T$, and $f_e(x) > r$ for $e \in T'$, which is the same as saying $\pi_f((x, r)) = \langle \emptyset, T, T' \rangle$, which in turn holds if and only if $(x, r) \in C_{\langle \emptyset, T, T' \rangle}$. \square

Theorem 3.1 and Lemma 2.1 immediately yield the important

Corollary 3.1. *If f_E and g_E are order-equivalent function collections, then $\text{VOD}(f_E) = \text{VOD}(g_E)$.*

Applying Theorem 3.1 to collections of affine functions or collections of quadratic functions on \mathbf{R}^d and using the algorithmic facts stated in Section 2 we obtain the following important unifying algorithmic result.

Theorem 3.2. *Let $D = \mathbf{R}^d$ for some $d \geq 1$. Let f_E be a collection of affine functions on \mathbf{R}^d given by*

$$f_e(x) = \langle a_e, x \rangle + b_e, \quad a_e \in \mathbf{R}^d, b_e \in \mathbf{R}$$

or let f_E be a collection of quadratic functions of the form

$$f_e(x) = (x - p_e)^T A(x - p_e) + t_e, \quad p_e \in \mathbf{R}^d, t_e \in \mathbf{R},$$

where A is a $d \times d$ real symmetric matrix.

$\text{VOD}(f_E)$ can be constructed in worst case time

$$O(n \log n) \quad \text{for } d = 1, 2 \text{ and}$$

$$O(n^{\lceil (d+1)/2 \rceil}) \quad \text{for } d \geq 3.$$

This is optimal for odd d and for $d = 2$.

Proof. By Theorem 3.1 and Corollary 3.1 it suffices to construct the one-level cells in an arrangement of hyperplanes in \mathbf{R}^{d+1} , which can be done in the given time bounds by Fact 2.4. \square

Theorem 3.2 is quite important because it unifies a number of algorithmic results in the literature about different kinds of VoDs. Below we give examples of what kind of planar VoDs are generated by collections of quadratic (or, equivalently, affine) functions.

Example 3.1. (Ordinary Euclidean VoD [31].) For each $e \in E$ let p_e be a point in \mathbf{R}^2 . Let I denote the 2×2 identity matrix. For each $e \in E$ let g_e be the quadratic function

$$g_e(x) = (x - p_e)^T I(x - p_e).$$

Note that $g_e(x)$ is the square of the Euclidean distance function between p_e and x , i.e.,

$$g_e(x) = (f_e(x))^2, \quad \text{where } f_e(x) = d(p_e, x).$$

Earlier we argued that $\text{VOD}(f_E)$ is the traditional Euclidean Voronoi diagram of the points p_e . As the distance function is nonnegative, the collection g_E is order-equivalent to f_E , and hence $\text{VOD}(g_E)$ is the traditional Euclidean Voronoi diagram as well.

Note 3.1. The ordinary Euclidean VoD is of such importance that it seems worthwhile to spell out again the geometric intuition that is hidden behind the formalism.

For each $e \in E$ the function g_e describes a paraboloid of rotation g_e^0 in \mathbb{R}^3 which is tangent to the x_1 - x_2 -plane at point p_e , has its axis parallel to the z -axis, and “opens upward” towards $z = +\infty$. Imagine the set $\{g_e^0 | e \in E\}$ of paraboloids in \mathbb{R}^3 penetrating each other. Furthermore, imagine each paraboloid to be opaque and having a unique color. Finally, imagine an observer standing at $z = -\infty$ looking in the positive z -direction. The visible parts of the paraboloids would appear like the VoD of the point set $\{p_e | e \in E\}$ with each V-region colored differently.

Now imagine that the entire 3-space is “warped” by the function F_I , with

$$F_I(x_1, x_2, z) = (x_1, x_2, z - x_1^2 - x_2^2).$$

Since the mapping F_I leaves the x_1 and x_2 coordinates invariant, and since it also does not change the difference in z -coordinate of points with identical x_1 - x_2 -coordinates, nothing changes for the observer at $z = -\infty$. He still sees the same colored tiling of the plane.

The point is, however, that F_I maps every paraboloid g_e^0 to the plane f_e^0 , where

$$f_e(x) = -2\langle p_e, x \rangle + \langle p_e, p_e \rangle.$$

Thus what the observer at $z = -\infty$ sees, is the projection onto the x_1 - x_2 -plane of the boundary of the polyhedron P formed by the intersection of the lower halfspaces $\{f_e^+ | e \in E\}$. Algorithmically, of course, this means that in order to construct the VoD it suffices to construct the polyhedron P . As a matter of fact, if one analyzes Hoey and Shamos' Voronoi diagram algorithm, it is indeed a disguised halfspace intersection algorithm, namely a dual version of Preparata and Hong's convex hull algorithm [28].

Finally, we want to point out that the planes f_e^0 are not arbitrary planes in \mathbb{R}^3 , but each f_e^0 is the tangent plane to the “upside down” paraboloid of rotation $q^0 = \{(x_1, x_2, z) | z = -x_1^2 - x_2^2\}$ at point p_e' which has the same x_1 and x_2 coordinates as p_e but has z -coordinate $-\langle p_e, p_e \rangle$.

Example 3.2. (Euclidean furthest point VoD [31].) Let the set E and the matrix I be as in Example 3.1. Again, for $e \in E$ let p_e be a point in the plane and let g_e be a function defined by

$$g_e(x) = (x - p_e)^T (-I)(x - p_e).$$

For the collection g_E of functions of this kind, $\text{VOD}(g_E)$ turns out to be what has been traditionally known as the furthest point Voronoi diagram of the planar point set $\{p_e | e \in E\}$.

Note 3.2. One of the reasons why the Euclidean closest point and furthest point Voronoi diagram are of such importance is the usefulness of their geometric dual graphs, the so-called closest point and furthest point *Delaunay triangulations* [31]. It seems worthwhile to point out that these triangulations can be naturally derived from the polyhedron P in the paraboloid construction in Note 3.1.

Let X be the set of points $\{p_e | e \in E\}$ on the upside down paraboloid q^0 (as in Note 3.1). Let P^* be the convex hull of X . Then the closest point Delaunay triangulation of the planar point set $\{p_e | e \in E\}$ is exactly the projection of the faces of P^* which are “on top of” P^* , i.e., which could be seen by an observer at $z = +\infty$ looking in the negative z -direction. Similarly, the furthest point Delaunay triangulation is exactly the projection of the faces of P^* which are “underneath”, i.e., the faces visible to an observer at $z = -\infty$.

Example 3.3. (Power diagrams [1], or VoDs in the “Laguerre geometry” [19].) Again let I denote the 2×2 identity matrix. For $e \in E$ let $p_e = (x_e, y_e)$ be a point in the plane and let t_e be a real number. Define functions g_e by

$$g_e(x) = (x - p_e)^T I (x - p_e) + t_e.$$

The VoD defined by the collection g_E of such functions was first discussed in the context of computational geometry in [1] where it is called “power diagram,” and in [19] where it is called Voronoi diagram in the Laguerre geometry. Apparently, it is also known as a Dirichlet cell complex [27].

Note 3.3. We want to point out that VoDs as defined in Example 3.3 model a seemingly natural growth process. Imagine a set of nonoverlapping circular cells C_e in the plane, each increasing with time in such a way that the growth rate of its radius is inversely proportional to its diameter. Whenever two cells come into contact, they cease to grow in the directions in which overlap would occur. The eventual shape of each cell of such a growth system is given by a VoD as in Example 3.3, where the collection g_E consists of a function g_e for each cell C_e , where p_e is the center of C_e , and t_e is the time when it started growing, i.e., the last time when its radius would have been zero.

This can easily be seen as follows: The growth rate of each cell C_e defines a real valued function g_e telling for each point in the plane at what time it would be covered by C_e if it grew uninhibitedly. For each point in the plane one wants to know which cell(s) overgrew it first. Thus the Voronoi diagram of the functions in g_E is indeed the desired object. Now it remains to show that the functions g_e have the form given in Example 3.3.

We want the radius of an uninhibited circular cell C_e with center p_e to grow at a rate inversely proportional to the diameter of C_e starting at time t_e . If $r(t)$ denotes the radius changing with time, then it must satisfy the differential equation

$$\frac{dr}{dt} = \frac{1}{2r(t)}$$

with initial condition $r(t_e) = 0$. The solution to this equation is given by $r(t) = \sqrt{t - t_e}$, or $r(t)^2 = t - t_e$. Expressing the radius $r(t)$ in terms of Cartesian coordinates yields

$$t = (x - p_e)^T I (x - p_e) + t_e.$$




Fig. 3.2.

Example 3.4. For $e \in E$ let p_e again be a point in the plane, and let J be a 2×2 diagonal matrix with entries 1 and -1 .

Let g_E be the collection of functions g_e , with

$$g_e(x) = (x - p_e)^T J(x - p_e).$$

The way $\text{VOD}(g_E)$ partitions the plane is rather peculiar. It appears to be a mixture of a Euclidean closest point VoD in the x -direction and an Euclidean furthest point VoD in the y -direction. An example is given in Fig. 3.2. To our knowledge this type of planar VoD has appeared in the literature only once, in a marginal remark in [19]. It is unclear whether it has any applications.

4. Order- k and Degree- k Voronoi Diagrams

Recall our definition of a Voronoi diagram at the beginning of the previous section. We defined it via an equivalence relation ρ_f on a domain D with respect to a set f_E of real valued functions on D , where ρ_f was defined by

$$x \rho_f y \Leftrightarrow R_f(x) = R_f(y),$$

where $R_f(x) = \{e \in E | f_e(x) = \min_{i \in E} f_i(x)\}$.

Obviously, there are ways of generalizing this definition by replacing the equivalence relation ρ_f by another one which is not based on the minimum of the function values, but rather on the k th smallest. With this in mind we give the following definition of a k -minimum: Let E be an index set, and for $i \in E$ let $x_i \in \mathbf{R}$. For an integer k , $k: \min_{i \in E} x_i$ is then defined to be the least real number z such that $x_i \leq z$ for k elements i of E .

We can now generalize the mapping R_f of Section 3 in two interesting ways. For an integer k , $1 \leq k \leq n$, define the mappings R_f^k and S_f^k from D to subsets of E by

$$R_f^k(x) = \left\{ e \in E | f_e(x) = k: \min_{i \in E} f_i(x) \right\} \quad \text{and}$$

$$S_f^k(x) = \left\{ e \in E | f_e(x) \leq k: \min_{i \in E} f_i(x) \right\}.$$

For every k , R_f^k and S_f^k induce equivalence relations ρ_f^k and σ_f^k on D ,

respectively, here for $x, y \in D$

$$x\rho_f^k y \Leftrightarrow R_f^k(x) = R_f^k(y) \quad \text{and}$$

$$x\sigma_f^k y \Leftrightarrow S_f^k(x) = S_f^k(y).$$

For $1 \leq k \leq n$ we call the partition of D induced by the equivalence relation ρ_f^k the *degree- k Voronoi diagram on D with respect to f_E* , or *degree- k -VOD(f_E)* for short. We call each equivalence class in the degree- k -VOD(f_E) a *degree- k V-cell*, and we denote such a cell by $k:V_T$, where $T \subseteq E$, and

$$k:V_T = \{x \in D \mid R_f^k(x) = T\}.$$

As in the case of the ordinary VoD, as defined in the previous section, we call a degree- k V-cell $k:V_T$ a *degree- k Voronoi region* if $|T| = 1$.

In the same manner we define the *order- k Voronoi diagram on D with respect to f_E* , *order- k -VOD(f_E)* for short, to be the partition of D induced by σ_f^k . We call each equivalence class an *order- k V-cell* and denote it by $k:W_T$, where

$$k:W_T = \{x \in D \mid S_f^k(x) = T\}.$$

Order- k Voronoi regions are order- k V-cells $k:W_T$ with $|T| = k$.

It should be clear that both order-1-VOD(f_E) and degree-1-VOD(f_E) are the same as VOD(f_E). In the Euclidean case (i.e., f_E is chosen as in Example 3.1) the degree- k Voronoi diagram partitions the plane by the k th nearest neighbor, whereas the order- k Voronoi diagram partitions the plane by the k nearest neighbors.

Both the order- k and the degree- k Voronoi diagram appear to be fairly natural generalizations of the ordinary VoD. So it is somewhat surprising that except for a short remark in [10] so far only the order- k VoD has been considered in the computational geometry literature [4], [11], [21], [31].¹ The main reason for this appears to be the fact that in the Euclidean case order- k V-regions are always convex and connected, whereas degree- k V-regions generally consist of several convex components.

As in the case of VOD(f_E), and perhaps even more so, it turns out to be beneficial to view degree- k -VOD(f_E) and order- k -VOD(f_E) as derived from the f_E -arrangement in $D \times \mathbf{R}$. To this end we have the following theorems.

Theorem 4.1. *For $e \in E$ let f_e be a real valued function on D . For some $T \subseteq E$ let $k:V_T$ be a degree- k V-cell:*

$$k:V_T = \bigcup \left\{ \text{proj}(C) \mid C \text{ is a } k\text{-level cell in the } f_E\text{-arrangement of the form } C = C_{\langle E^-, T, E^+ \rangle} \right\}.$$

¹The reader may convince himself that our definition of order- k -VOD(f_E) actually agrees with the definitions offered in these papers, when f_E is chosen as in Example 3.1 of the previous section.

Proof. Let $x \in k:V_T$ and let $r = k: \min_{i \in E} f_i(x)$. By definition we have $f_e(x) = r$ for $e \in T$. Let $E^- = \{e | f_e(x) < r\}$. By the definition of the $k: \min$ it must be the case that $|E^-| < k$, but $|T| + |E^-| \geq k$. But this is the case if and only if $(x, r) \in C_{\langle E^-, T, E^+ \rangle}$, where $E^+ = (E - T) - E^-$ and $C_{\langle E^-, T, E^+ \rangle}$ is a k -level cell. \square

We state similar theorems for the order- k Voronoi diagram. The proofs follow directly from the definition and are omitted.

Theorem 4.2. For some $T \subseteq E$ let $k:W_T$ be an order- k V-cell.

$$k:W_T = \bigcup \left\{ \text{proj}(C) \mid C \text{ is a } k\text{-level cell in the } f_E\text{-arrangement with} \right. \\ \left. C = C_{\langle E^-, E^0, E^+ \rangle} \text{ and } T = E^- \cup E^0 \right\}.$$

Even more useful might be the following

Theorem 4.3. For some $T \subseteq E$ let $k:W_T$ be an order- k V-cell.

If $k:W_T$ is an order- k Voronoi region, then

$$k:W_T = \text{proj}(C_{\langle T, \emptyset, E-T \rangle}) \text{ and} \\ C_{\langle T, \emptyset, E-T \rangle} \text{ is a } k\text{-belt cell in the } f_E\text{-arrangement.}$$

Otherwise

$$k:W_T = \bigcup \left\{ \text{proj}(C_{\langle E^-, E^0, E^+ \rangle}) \mid E^- \cup E^0 = T \text{ and } C_{\langle E^-, E^0, E^+ \rangle} \text{ is} \right. \\ \left. \text{a } k\text{-level cell as well as a } (k+1)\text{-level cell} \right\}.$$

Figures 4.1 and 4.2 illustrate the contents of the preceding theorems.

The preceding theorems and Lemma 2.1 have an important corollary.




Fig. 4.1. A degree-2 Voronoi Diagram.

5. Intersection Problems

In this section we want to show briefly how an f_E -arrangement can be a useful tool for solving certain intersection or union problems for finite sets of objects. Note that using de Morgan's laws, union can be reduced to intersection. We therefore restrict ourselves to intersection problems. At first we describe our method abstractly, later we apply it to solve intersection problems for finite sets of discs and, more generally, finite sets of similar conic sections in \mathbb{R}^d .

For $e \in E$ let B_e be some subset of our domain D and assume for each e there is a real valued function f_e on D with $f_e(x) < 0$ iff $x \in B_e$. We are interested in describing $\bigcap_{e \in E} B_e$ in terms of the f_E -arrangement. For this purpose let us identify every subset $B \subseteq D$ with its injection $\{(x, 0) | x \in B\}$ into $D \times \mathbb{R}$. This way we can write for every $e \in E$ that $B_e = f_e^- \cap D$.

The intersection $\bigcap_{e \in E} B_e$ can then be written as

$$\bigcap_{e \in E} (f_e^- \cap D) = D \cap \bigcap_{e \in E} f_e^- ,$$

which is of course nothing but $D \cap C_{\langle E, \emptyset, \emptyset \rangle}$.

This means one can construct $\bigcap_{e \in E} B_e$ by first constructing the "top" full cell $C_{\langle E, \emptyset, \emptyset \rangle}$ of the f_E -arrangement and then intersecting this cell with the "base plane" $\{(x, 0) | x \in D\}$.

As an example consider the problem of constructing the intersection of a finite set of open discs in the Euclidean plane \mathbb{R}^2 , where each disc B_e has center $c_e \in \mathbb{R}^2$ and radius $r_e > 0$.

Associate with each disc B_e the real valued function f_e on \mathbb{R}^2 , where for $x \in \mathbb{R}^2$

$$f_e(x) = (x - c_e)^T I (x - c_e) - r_e^2 ,$$

and I is the identity matrix.

Note that as desired the condition $x \in B_e$ iff $f_e(x) < 0$ holds. Therefore the idea outlined above is applicable.

First construct the cell $C_{\langle E, \emptyset, \emptyset \rangle}$ formed by the intersection of the upper hemispaces f_e^- . As all the functions f_e are quadratic functions generated by the same matrix I , the results of Section 2 imply that $C_{\langle E, \emptyset, \emptyset \rangle}$ can be constructed in $O(n \log n)$ time. This cell can then be intersected with the base plane $\{(x, 0) | x \in \mathbb{R}^2\}$ in $O(n)$ time to yield the intersection $\bigcap_{e \in E} B_e$ of all the discs in $O(n \log n)$ time overall. The same approach can be used to construct the intersection of n open balls in \mathbb{R}^d in time $O(n \lceil (d+1)/2 \rceil)$ for $d > 2$.

This method of constructing the intersection of discs is very closely related to the one proposed by Brown [8]. He uses spheres and spherical inversion to reduce the problem to one of the intersecting halfspaces, whereas we, if one analyzes our method in detail, use paraboloids and the "warping" function F_f given at the end of Section 2.

We briefly want to mention some more general applications of our method of using the f_E -arrangement formalism to solve intersection problems. One concerns

intersecting regions B_e of the form

$$B_e = \{x \in \mathbf{R}^d \mid f_e(x) < 0\},$$

where

$$f_e(x) = (x - c_e)^T A (x - c_e) + r_e, \quad c_e \in \mathbf{R}^d, r_e \in \mathbf{R},$$

and A a fixed nonsingular symmetric matrix.

Such regions B_e are general conic sections. For instance, in the case $d = 2$ and A indefinite, B_e might be the unbounded (nonconvex) region between the two branches of a hyperbola. With our formalism the intersection of n such regions B_e can be accomplished in $O(n \log n)$ time in case $d = 2$, and $O(n^{\lceil (d+1)/2 \rceil})$ time in case $d > 2$. Note that Brown's method cannot be used for this purpose as it relies on spherical inversion.

The machinery of f_E -arrangements can also be used to solve intersection problems of the form "find all x that lie in exactly k regions B_e ." Finding such a set reduces to the problem of intersecting the k -belt cells of an f_E -arrangement with the "base plane." Similarly, the boundary of the set of all x that lie in at least k regions B_e can be constructed via intersecting the k -level cells with the "base plane". We leave the details to the reader.

6. Conclusions

The initial seed for the ideas in this paper was the observation that Voronoi diagrams of point sets in the plane are related to three-dimensional polyhedra whose facets are tangent to a common paraboloid. We discovered this while scrutinizing Brown's use of spherical inversion to relate Voronoi diagrams in the plane with three-dimensional polytopes whose vertices all lie on a common sphere [7].² At first these relationships seemed rather mysterious and inexplicable. Only after we turned our attention to the question of what Voronoi diagrams really are, did we arrive at the (we think) satisfying explanation of these relationships presented in this paper.

An important concept in this paper is the notion of an f_E -arrangement over a domain D , i.e., the partition of $D \times \mathbf{R}$ induced by a finite collection f_E of real valued functions on D . Of course, in its full generality this concept is quite useless. However, with appropriate restrictions f_E -arrangements can provide interesting geometric-combinatorial research. For instance, if $D = \mathbf{R}^d$ and the functions in f_E are continuous and satisfy certain simple finite intersection properties, f_E -arrangements have been studied fairly extensively as *pseudoplane arrangements* [17], [18] and also in the context of *oriented matroids* [26]. From the point of view of computational geometry, pseudoplane arrangements still offer

²Actually, it turns out that the use of paraboloids and spheres to relate Voronoi diagrams and polyhedra are projectively equivalent.

some interesting algorithmic problems. How difficult is it, for instance, to construct all cells in such an arrangement when only f_E is given? It seems that the $O(n^{d+1})$ algorithm in [14] actually generalizes to that case. How difficult is it to construct all k -level cells? In particular, for $D = \mathbb{R}^2$ is it possible to construct the one-level cells in $O(n \log n)$ time? This would yield a fast construction algorithm for a large class of planar Voronoi diagrams.

Finally there is the question whether there are other interesting classes of f_E -arrangements. There appear to be several natural and promising ways of arriving at such arrangements. One would be to relax or change the finite intersection properties of functions postulated in the case of pseudoplane arrangements. For instance, in the case $D = \mathbb{R}^2$, allowing two surfaces to intersect and cross either in a line or in a simple closed curve would give rise to a class of arrangements that includes the ones that correspond to weighted Voronoi diagrams [3]. Another interesting way of generalizing would be to change the underlying domain D to, say, a torus, and to try to postulate appropriate intersection properties for that case.

Acknowledgments

We would like to thank John Gilbert for his careful reading of the manuscript and his many suggestions for improvement. We also want to thank Bennett Battaile, Gianfranco Bilardi, Joseph O'Rourke, and Chee Yap for their comments.

References

1. F. Aurenhammer, Power diagrams: Properties, algorithms and applications, Rep. F120, IIG, Tech. Univ. Graz, Austria, 1983.
2. F. P. Ash and E. D. Boler, Generalized Dirichlet tessellations, to be published.
3. F. Aurenhammer and H. Edelsbrunner, An optimal algorithm for constructing the weighted Voronoi diagram in the plane, *Pattern Recognition* 17 (1984), 251–257.
4. B. Bhattacharya, An algorithm for computing order k Voronoi diagrams in the plane, CS Dept. Tech. Rep. TR. 83-9, Simon Fraser Univ., Burnaby B.C., 1983.
5. A. Bowyer, Computing Dirichlet tessellations, *The Computer Journal* 24 (1981), 162–166.
6. K. Q. Brown, Fast intersection of half spaces, CS Dept. Rep. CMU-CS-78-129, CMU, Pittsburgh PA, 1978.
7. K. Q. Brown, Voronoi diagrams from convex hulls, *Info. Proc. Let.* 9 (1979), 223–228.
8. K. Q. Brown, Geometric transforms for fast geometric algorithms, Ph.D. thesis, CS Dept. Rep. CMU-CS-80-101, CMU, Pittsburgh PA, 1980.
9. L. P. Chew and R. L. Drysdale, III, Voronoi diagrams based on convex distance functions, *Proc. of the ACM Symp. on Computational Geometry*, 1985, 235–244.
10. R. Cole and C. K. Yap, Geometric retrieval problems, *Proc. of the 24th IEEE Symp. on Foundations of Computer Science*, 1983, 112–121.
11. F. Dehne, An optimal algorithm to construct all Voronoi diagrams for k nearest neighbor searching in the Euclidean plane, *Proc. of the 20th Annual Allerton Conf. on Communication, Control, and Computing*, 1982.
12. R. L. Drysdale, III, Generalized Voronoi diagrams and geometric searching, Ph.D. thesis, CS Dept. Rep. STAN-CS-79-705, Stanford Univ., Stanford, CA, 1979.
13. H. Edelsbrunner, *Arrangements and Geometric Computations*, forthcoming book.

14. H. Edelsbrunner, J. O'Rourke, and R. Seidel, Constructing arrangements of lines and hyperplanes with applications, Proc. 24th Symp. on Foundations of Computer Science, 1983, 83–91 (to appear in SIAM J. Computing).
15. H. Edelsbrunner and E. Welzl, On the number of line-separations of a finite set in the plane, Rep. F97, IIG, Tech. Univ. Graz, Austria, 1982 (to appear in J. Combin. Theory, Ser. A).
16. P. Erdős, L. Lovasz, A. Simmons, and E. G. Straus, Dissection graphs of planar point sets, in A Survey of Combinatorial Theory, J. N. Srivastava et al., eds., North-Holland, Amsterdam, 1973.
17. B. Grünbaum, Arrangements of hyperplanes, Conf. Numerantium III, Louisiana Conf. on Comb., Graph Theory and Computing, 1971, 41–106.
18. B. Grünbaum, Arrangements and spreads, Reg. Conf. Series in Math., AMS, Providence, 1972.
19. H. Imai, M. Iri, and K. Murota, Voronoi diagrams in the Laguerre geometry and its application, SIAM J. Computing, to appear.
20. D. G. Kirkpatrick, Efficient computation of continuous skeletons, Proc. of the 20th IEEE Symp. on Foundations of Computer Science, 1979, 18–27.
21. D. T. Lee, On finding k -nearest neighbors in the plane, M.S. thesis, Coordinated Science Lab., Rep. R-728, Univ. of Illinois, Urbana Ill., 1976.
22. D. T. Lee, Two dimensional Voronoi diagrams in the L_p -metric, J. ACM, Oct. 604–618 (1980).
23. D. T. Lee and R. L. Drysdale, III, Generalized Voronoi diagrams in the plane, SIAM J. Computing, 10 (1981), 73–87.
24. D. T. Lee and C. K. Wong, Voronoi diagrams in L_1 - (L_∞ -) metrics with 2-dimensional storage applications, SIAM J. Computing, 9 (1980), 200–211.
25. P. McMullen, The maximum number of faces of a convex polytope, Mathematika 17 (1970), 179–184.
26. A. Mandel, Topology of oriented matroids, Ph.D. thesis, Dept. of Combinatorics and Optimization, Univ. of Waterloo, Waterloo, Ont., 1981.
27. I. Paschinger, Konvexe Polytope und Dirichletsche Zellenkomplexe, Math. Inst., Univ. Salzburg, Austria, Arbeitsbericht 1–2, 1982.
28. F. P. Preparata and S. J. Hong, Convex hulls of finite sets of points in two and three dimensions, C. ACM 20 (1977), 87–93.
29. F. P. Preparata and D. E. Muller, Finding the intersection of n halfspaces in time $O(n \log n)$, Theore. Comput. Sci. 8 (1979), 45–55.
30. R. Seidel, A convex hull algorithm optimal for point sets in even dimensions, Rep. 81-14, Dept. of CS, Univ. of BC, Vancouver B.C., 1981.
31. M. I. Shamos and D. Hoey, Closest-point problems, Proc. of the 17th IEEE Symp. on Foundations of Computer Science, 208–215, 1975.
32. M. Sharir, and D. Leven, Intersection problems and application of Voronoi diagrams, in Advances in Robotics, Vol. 1: Algorithmic and Geometric Aspects of Robotics, J. Schwartz and C. K. Yap, eds., Lawrence Erlbaum Associates Inc. (to appear).
33. D. F. Watson, Computing the n -dimensional Delaunay triangulation with applications to Voronoi polytopes, Comput. J. 24 (1981), 167–172.
34. C. K. Yap, An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments, manuscript, Courant Institute of Math. Sciences, NYU, New York, NY, 1984.
35. T. Zaslavsky, Facing up to arrangements: Face-count formulas for partitions of space by hyperplanes, Memoirs AMS 154 (1975).

Received March 29, 1985, and in revised form October 1, 1985.