2

Curved Voronoi Diagrams

Jean-Daniel Boissonnat*, Camille Wormser, and Mariette Yvinec

Abstract

Voronoi diagrams are fundamental data structures that have been extensively studied in Computational Geometry. A Voronoi diagram can be defined as the minimization diagram of a finite set of continuous functions. Usually, each of those functions is interpreted as the distance function to an object. The associated Voronoi diagram subdivides the embedding space into regions, each region consisting of the points that are closer to a given object than to the others. We may define many variants of Voronoi diagrams depending on the class of objects, the distance functions and the embedding space. Affine diagrams, i.e. diagrams whose cells are convex polytopes, are well understood. Their properties can be deduced from the properties of polytopes and they can be constructed efficiently. The situation is very different for Voronoi diagrams with curved regions. Curved Voronoi diagrams arise in various contexts where the objects are not punctual or the distance is not the Euclidean distance. We survey the main results on curved Voronoi diagrams. We describe in some detail two general mechanisms to obtain effective algorithms for some classes of curved Voronoi diagrams. The first one consists in linearizing the diagram and applies, in particular, to diagrams whose bisectors are algebraic hypersurfaces. The second one is a randomized incremental paradigm that can construct affine and several planar non-affine diagrams.

We finally introduce the concept of Medial Axis which generalizes the concept of Voronoi diagram to infinite sets. Interestingly, it is possible to efficiently construct a certified approximation of the medial axis of a bounded set from the Voronoi diagram of a sample of points on the boundary of the set.

* Chapter coordinator

2.1 Introduction

Voronoi diagrams are fundamental data structures that have been extensively studied in Computational Geometry. Given n objects, the associated Voronoi diagram subdivides $\mathbb{R}^d$ into regions, each region consisting of the points that are closer to a given object than to any other object. We may define many variants of Voronoi diagrams depending on the class of objects, the distance function and the embedding space. Although Voronoi diagrams are most often defined in a metric setting, they can be defined in a more abstract way. In Sect. 2.2, we define them as minimization diagrams of any finite set of continuous functions without referring to a set of objects.

Given a finite set of objects and associated distance functions, we call bisector the locus of the points that are at equal distance from two objects. Voronoi diagrams can be classified according to the nature of the bisectors of the pairs of objects, called the bisectors of the diagram for short. An important class of Voronoi diagrams is the class of affine diagrams, whose bisectors are hyperplanes. Euclidean Voronoi diagrams of finite point sets are affine diagrams. Other examples of affine diagrams are the so-called power (or Laguerre) diagrams, where the objects are no longer points but hyperspheres and the Euclidean distance is replaced by the power of a point to a hypersphere. In Sect. 2.3, we recall well-known facts about affine diagrams. In particular, we characterize affine diagrams and establish a connection between affine diagrams and polytopes. As a consequence, we obtain tight combinatorial bounds and efficient algorithms. We also obtain a dual structure that is a triangulation under a general position assumption.

Non-affine diagrams are by far less well understood. Non-affine diagrams are obtained if one changes the distance function: additively and multiplicatively weighted distances are typical examples. Such diagrams allow to model growing processes and have important applications in biology, ecology, chemistry and other fields (see Sect. 2.9). Euclidean Voronoi diagrams of non-punctual objects are also non-affine diagrams. They are of particular interest in robotics, CAD and molecular biology. Even for the simplest diagrams, e.g. Euclidean Voronoi diagrams of lines, triangles or spheres in 3-space, obtaining tight combinatorial bounds, efficient algorithms and effective implementations are difficult research questions.

A first class of non-affine diagrams to be discussed in Sect. 2.4 is the case of diagrams whose bisectors are algebraic hypersurfaces. We first consider the case of Möbius diagrams whose bisectors are hyperspheres and the case of anisotropic diagrams whose bisectors are quadratic hypersurfaces (see Sect. 2.4.2). The related case of Apollonius (or Johnson-Mehl) diagrams is also described in Sect. 2.4.

The key to obtaining effective algorithms for computing those non-affine diagrams is a linearization procedure that reduces the construction of a non-affine diagram to intersecting an affine diagram with a manifold in some higher dimensional space. This mechanism is studied in full generality in Sect. 2.5.

In this section, we introduce abstract diagrams, which are diagrams defined in terms of their bisectors. By imposing suitable conditions on these bisectors, any abstract diagram can be built as the minimization diagram of some distance functions, thus showing that the class of abstract diagrams is the same as the class of Voronoi diagrams. Furthermore, the linearization technique introduced in Sect. 2.5 allows to prove that if the bisectors of a diagram belong to a certain class of bisectors, the distance functions defining the diagram can be chosen among a precise class of functions. For instance, affine diagrams are identified with power diagrams, spherical diagrams are identified with Möbius diagrams, and quadratic diagrams with anisotropic diagrams.

In Sect. 2.6, we introduce the incremental paradigm for constructing various diagrams. Under some topological conditions to be satisfied by the diagram, the incremental construction is efficient. The algorithm can be further improved by using a randomized data structure called the Voronoi hierarchy that allows fast localization of new objects. We then obtain fast randomized incremental algorithms for affine diagrams in any dimension and several non-affine diagrams in the plane. Going beyond those simple cases is difficult. As mentioned above, tight combinatorial bounds and efficient algorithms are lacking even for simple cases. Moreover, the numerical issues are delicate and robust implementations are still far ahead of the state of the art. This motivates the quest for approximate solutions.

In Sect. 2.7, we introduce the concept of Medial Axis of a bounded set Ω, which can be seen as an extension of the notion of Voronoi diagram to infinite sets. Interestingly, it is possible to construct certified approximations of the medial axis of quite general sets efficiently. One approach to be described consists in sampling the boundary of Ω and then computing an appropriate subset of the Voronoi diagram of the sample that approximates the medial axis. Hence the problem of approximating the medial axis of Ω boils down to sampling the boundary of Ω, a problem that is closely related to mesh generation (see Chap. 5).

Sect. 2.8 is devoted to the main CGAL software packages for computing Voronoi diagrams. Sect. 2.9 discusses some applications of curved Voronoi diagrams.

This chapter focuses on curved Voronoi diagrams defined in $\mathbb{R}^d$ and aims at providing useful background and effective algorithms. Additional material can be found in surveys on Voronoi diagrams [276, 37] and in text books on Computational Geometry [110, 67]. This chapter does not consider Voronoi diagrams defined in more general spaces. Voronoi diagrams can be defined in hyperbolic geometry without much difficulty [60, 67]. In the Poincaré model of hyperbolic geometry, the bisectors are hyperspheres and hyperbolic diagrams of finite point sets are a special case of Möbius diagrams. Computing Voronoi diagrams on Riemannian manifolds is much more involved and very few is known about such diagrams and their construction [239].

Notation: We identify a point $x \in \mathbb{R}^d$ and the vector of its coordinates. We note $x \cdot y$ the dot product of x and y, $x^2 = x \cdot x = \|x\|^2$ the squared Euclidean norm of x, and $\|x - y\|$ the Euclidean distance between points x and y.

We call *hypersurface* a manifold of codimension 1. Examples to be used in this chapter are hyperplanes, hyperspheres and quadratic hypersurfaces.

2.2 Lower Envelopes and Minimization Diagrams

Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a set of d-variate continuous functions defined over $\mathbb{R}^d$. The *lower envelope* of $\mathcal{F}$ is defined as

$$\mathcal{F}^- = \min_{1 \le i \le n} f_i.$$

From $\mathcal{F}$ and $\mathcal{F}^-$, we define a natural partition of $\mathbb{R}^d$ called the *minimization diagram* of $\mathcal{F}$. For a point $x \in \mathbb{R}^d$, we define the *index set* $I(x)$ of x as the set of all indices i such that $\mathcal{F}^-(x) = f_i(x)$. An equivalence relation noted $\equiv$ can then be defined between two points of $\mathbb{R}^d$ if they have the same index set:

$$x \equiv y \quad \Leftrightarrow \quad I(x) = I(y).$$

The equivalence classes $\mathbb{R}^d / \equiv$ are relatively open sets that cover $\mathbb{R}^d$. Their closures are called the *faces* of the minimization diagram of $\mathcal{F}$ (see Fig. 2.1). The index set of a face is defined as the largest subset of indices common to all the points of the face. Conversely, the face of index set I is the set of all points x such that $I \subset I(x)$.

Observe that the faces of this diagram are not necessarily contractible nor even connected. In particular, a 0-dimensional face may consist of several distinct points.

Lower envelopes and minimization diagrams have been well studied. We recall an important result due to Sharir [311] which provides an almost optimal result when the f_i are supposed to be multivariate polynomials of constant maximum degree.

Theorem 1 (Sharir). *The number of faces of the minimization diagram of a set $\mathcal{F}$ of n multivariate polynomials of constant maximum degree η is $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$, where the constant of proportionality depends on ε, d and η. The vertices, edges and 2-faces of the diagram can be computed in randomized expected time $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$.*

This general result is close to optimal in the worst-case (see Exercise 2). It has been improved in some special cases. For more information and other related results, one should consult the book by Sharir and Agarwal [312].

Voronoi diagrams, in their general setting, are just minimization diagrams of a finite set of continuous functions. This general definition encompasses

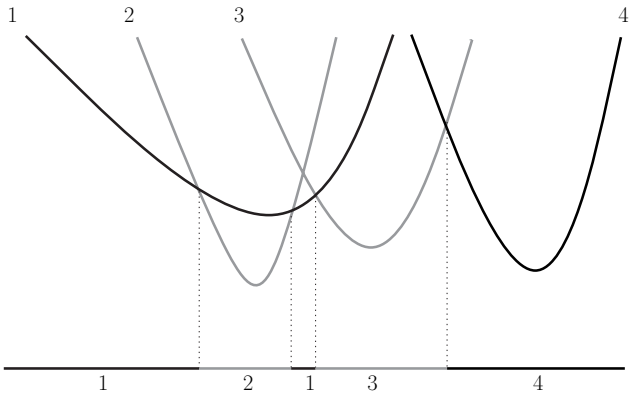

Fig. 2.1. The lower envelope of a set of univariate functions. The minimization diagram is drawn on the horizontal line with the corresponding indices. The face of index {1} consists of two components

the more traditional definition of Voronoi diagrams where the functions are defined as distance functions to a finite set of objects. Consider a set of objects $\mathcal{O} = \{o_1, \ldots, o_n\}$. To each object o_i is attached a continuous function δ_i that measures the distance from a point x of $\mathbb{R}^d$ to o_i. In the simplest case, $\mathcal{O}$ is a finite set of points and $\delta_i(x)$ is the Euclidean distance from x to o_i. The Voronoi diagram of $\mathcal{O}$ is defined as the minimization diagram of $\Delta = \{\delta_1, \ldots, \delta_n\}$. The concept of Voronoi diagram has been generalized and various other diagrams have been defined by considering more general objects and other distance functions. Distance is then not to be taken with too much rigor. The function δ_i is only supposed to be continuous.

Theorem 1 provides very general bounds on the complexity of Voronoi diagrams. However, this result calls for improvement. First, in some special cases, much better bounds can be obtained by other approaches to be discussed later in this chapter. In particular, we will see that the most popular Euclidean Voronoi diagram of points has a much smaller combinatorial complexity than the one given in the theorem.

A second issue is the algorithmic complexity. The algorithm mentioned in the theorem fails to provide a complete description of the diagram since only faces of dimensions up to 2 are computed.

Moreover, the implementation of such an algorithm remains a critical issue. As evidenced in Chap. 1, computing lower envelopes of algebraic functions is a formidable task, even in the simplest cases, e.g. quadratic bi-variate functions. We do not know of any implementation for higher degrees and dimensions.

The main goal of the following sections is to present effective algorithms for a variety of Voronoi diagrams for which some additional structure can be exhibited.

Exercise 1. Show that the combinatorial complexity of the lower envelope of n univariate functions whose graphs intersect pairwise in at most two points is $O(n)$. Show that the envelope can be computed in optimal time $\Theta(n \log n)$.

Exercise 2. Show that the convex hull of n ellipsoids of $\mathbb{R}^d$ may have $\Omega(n^{d-1})$ faces. Since the non-bounded faces of the Euclidean Voronoi diagram of n objects are in 1-1 correspondence with the faces of their convex hull, we get a lower bound on the size of the Voronoi diagram of n ellipsoids of $\mathbb{R}^d$. (Hint: consider n ellipsoids inscribed in a $(d-1)$-sphere S and intersecting S along great n $(d-2)$-spheres $\sigma_1, \ldots, \sigma_n$. The arrangement of the σ_i has $\Theta(n^{d-1})$ faces.)

2.3 Affine Voronoi Diagrams

We first introduce Euclidean Voronoi diagrams of points and establish a correspondence between those diagrams and convex polyhedra in one dimension higher. Polarity allows to associate to a Voronoi diagram its dual cell complex, called a Delaunay triangulation.

Almost identical results can be obtained for power (or Laguerre) diagrams where points are replaced by hyperspheres and the Euclidean distance by the power of a point to a hypersphere. Power diagrams constitute a natural extension of Euclidean Voronoi diagrams and are still affine diagrams. In fact, we will see that any affine diagram is the power diagram of a finite set of hyperspheres.

2.3.1 Euclidean Voronoi Diagrams of Points

Let $\mathcal{P} = \{p_1, \ldots, p_n\}$ be a set of points of $\mathbb{R}^d$. To each p_i, we associate its Voronoi region $V(p_i)$

$$V(p_i) = \{x \in \mathbb{R}^d : \|x - p_i\| \leq \|x - p_j\|, \forall j \leq n\}.$$

The region $V(p_i)$ is the intersection of $n-1$ half-spaces. Each such half-space contains p_i and is bounded by the bisector of p_i and some other point of $\mathcal{P}$. Since the bisectors are hyperplanes, $V(p_i)$ is a convex polyhedron, possibly unbounded.

The *Euclidean Voronoi diagram* of $\mathcal{P}$, noted $\mathrm{Vor}(\mathcal{P})$, is the cell complex whose cells are the Voronoi regions and their faces. Equivalently, the Euclidean Voronoi diagram of $\mathcal{P}$ can be defined as the minimization diagram of the distance functions $\delta_i, \ldots, \delta_n$, where

$$\delta_i(x) = \|x - p_i\|.$$

In other words, the Euclidean Voronoi diagram of $\mathcal{P}$ is the minimization diagram of a set of functions whose graphs are vertical[1] cones of revolution of

[1] By vertical, we mean that the axis of revolution is perpendicular to $\mathbb{R}^d$.

$\mathbb{R}^{d+1}$. Since minimizing $\|x-p_i\|$ over i is the same as minimizing $(x-p_i)^2$, the Euclidean Voronoi diagram of $\mathcal{P}$ can alternatively be defined as the mimization diagram of the smooth functions $(x-p_i)^2$ whose graphs are translated copies of a vertical paraboloid of revolution of $\mathbb{R}^{d+1}$.

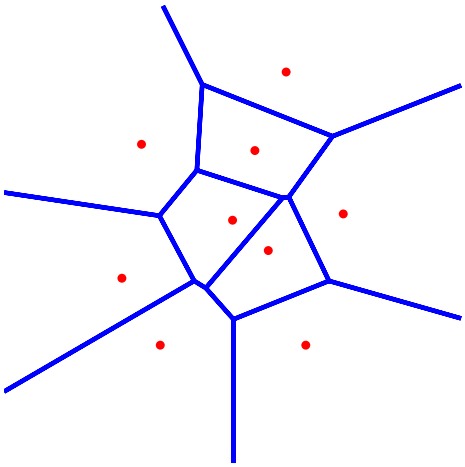

Fig. 2.2. The Voronoi diagram of a set of 9 points

Observing further that, for any x, $\arg\min_i(x-p_i)^2 = \arg\min_i(-2p_i \cdot x + p_i^2)$, we obtain that the Euclidean Voronoi diagram of $\mathcal{P}$ is the minimization diagram of a set of affine functions, namely the functions

$$d_i(x) = -2p_i \cdot x + p_i^2$$

whose graphs are hyperplanes of $\mathbb{R}^{d+1}$. Let us call h_{p_i}, $i = 1, \ldots, n$, those hyperplanes and let $h^-_{p_i}$ denote the half-space lying below h_{p_i}. The minimization diagram of the d_i is obtained by projecting the polyhedron

$$\mathcal{V}(\mathcal{P}) = h^-_{p_1} \cap \cdots \cap h^-_{p_n}.$$

vertically onto $\mathbb{R}^d$. See Fig. 2.3.

We have therefore proved the following theorem:

Theorem 2. *The faces of the Euclidean Voronoi diagram* $\mathrm{Vor}(\mathcal{P})$ *of a set of points* $\mathcal{P}$ *are the vertical projections of the faces of the convex polyhedron* $\mathcal{V}(\mathcal{P})$.

Exercise 3. Consider the maximization diagram obtained by projecting the faces of $h^+_{p_1} \cap \cdots \cap h^+_{p_n}$ vertically. Characterize the points that belong to a face of this diagram in terms of the distance to the points of $\mathcal{P}$.

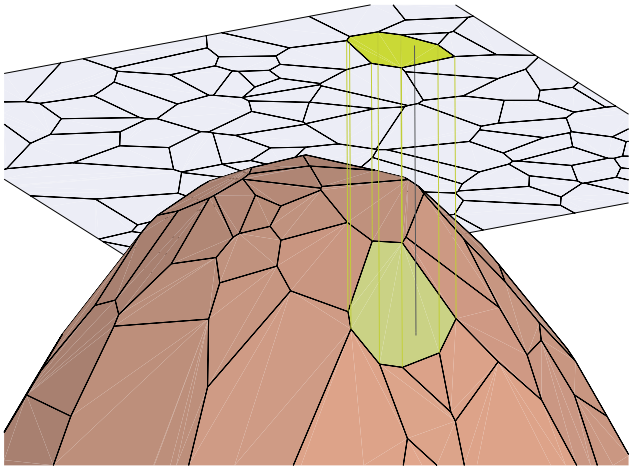

Fig. 2.3. The polyhedron $\mathcal{V}(\mathcal{P})$, with one of its faces projected onto $\mathbb{R}^d$

2.3.2 Delaunay Triangulation

Two cell complexes V and D are said to be dual if there exists an involutive correspondence between the faces of V and the faces of D that reverses the inclusions, i.e. for any two faces f and g of V, their dual faces f^* and g^* satisfy: $f \subset g \Rightarrow g^* \subset f^*$. We introduce now a cell complex that is *dual* to the Voronoi diagram of a finite set of points $\mathcal{P}$.

We assume for now that the set of points $\mathcal{P}$ is in *general position*, which means that no subset of $d+2$ points of $\mathcal{P}$ lie on a same hypersphere. Let f be a face of dimension k of the Voronoi diagram of $\mathcal{P}$. All points in the interior of f have the same subset $\mathcal{P}_f$ of closest points in $\mathcal{P}$. The face dual to f is the convex hull of $\mathcal{P}_f$. The *Delaunay triangulation of* $\mathcal{P}$, noted $\mathrm{Del}(\mathcal{P})$, is the cell complex consisting of all the dual faces. Because points of $\mathcal{P}$ are assumed to be in general position, $|\mathcal{P}_f| = d - k + 1$, all the faces of $\mathrm{Del}(\mathcal{P})$ are simplices and $\mathrm{Del}(\mathcal{P})$ is a simplicial complex. The fact that $\mathrm{Del}(\mathcal{P})$ is indeed a triangulation, i.e. a simplicial complex embedded in $\mathbb{R}^d$ and covering the convex hull of $\mathcal{P}$, will be proved now using a duality between points and hyperplanes in the so-called space of spheres.

Polarity

Let σ be the hypersphere of $\mathbb{R}^d$ of equation

$$\sigma(x) = (x - c)^2 - r^2 = x^2 - 2c \cdot x + s = 0,$$

where c is the center of σ, r its radius and $s = \sigma(0) = c^2 - r^2$.

We define the following bijective mapping

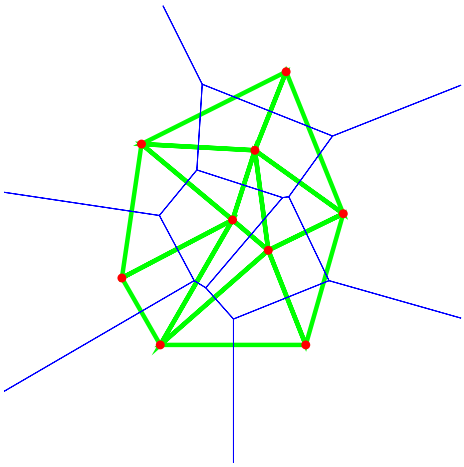

Fig. 2.4. The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines)

$$\phi : \sigma \in \mathbb{R}^d \longrightarrow \phi(\sigma) = (c, -s) \in \mathbb{R}^{d+1}$$

that maps a hypersphere of $\mathbb{R}^d$ to a point of $\mathbb{R}^{d+1}$. We thus consider $\mathbb{R}^{d+1}$ as the images by ϕ of the hyperspheres of $\mathbb{R}^d$ and call $\mathbb{R}^{d+1}$ *the space of spheres*. We note $\phi(p) = (p, -p^2)$ the image by ϕ of a point, considered as a hypersphere of radius 0. Observe that $\phi(p)$ is a point of the paraboloid $\mathcal{Q}$ of $\mathbb{R}^{d+1}$ of equation $x^2 + x_{d+1} = 0$. The points of $\mathbb{R}^{d+1}$ that lie above $\mathcal{Q}$ are images of imaginary hyperspheres whose squared radii are negative. The points below $\mathcal{Q}$ are images of real hyperspheres.

We now introduced a mapping between points and hyperplanes of the space of spheres, known as *polarity*. Polarity associates to the point $\phi(\sigma)$ its *polar hyperplane* h_σ which is the hyperplane of $\mathbb{R}^{d+1}$ of equation $2c \cdot x + x_{d+1} - s = 0$. Observe that the intersection of h_σ with $\mathcal{Q}$ projects vertically onto σ, and that h_σ is the affine hull of the image by ϕ of the points of σ. If p is a point of $\mathbb{R}^d$, the polar hyperplane h_p of $\phi(p)$ is the hyperplane tangent to $\mathcal{Q}$ at $\phi(p)$.

We deduce the remarkable following property: $x \in \sigma$ if and only if $\phi(x) = (x, -x^2) \in h_\sigma$ and σ encloses x if and only if $\phi(x) \in h_\sigma^+$, where h_σ^+ (resp. h_σ^-) denotes the closed half-space above (resp. below) h_σ. Indeed

$$\sigma(x) = 0 \Longleftrightarrow x^2 - 2c \cdot x + s = 0 \Longleftrightarrow \phi(x) \in h_\sigma$$
$$\sigma(x) < 0 \Longleftrightarrow x^2 - 2c \cdot x + s < 0 \Longleftrightarrow \phi(x) \in \operatorname{int} h_\sigma^+,$$

where $\operatorname{int} h_\sigma^+$ denotes the open half-space above h_σ.

Polarity is an involution that preserves incidences and reverses inclusions. Indeed, if σ and σ' are two hyperspheres, we have

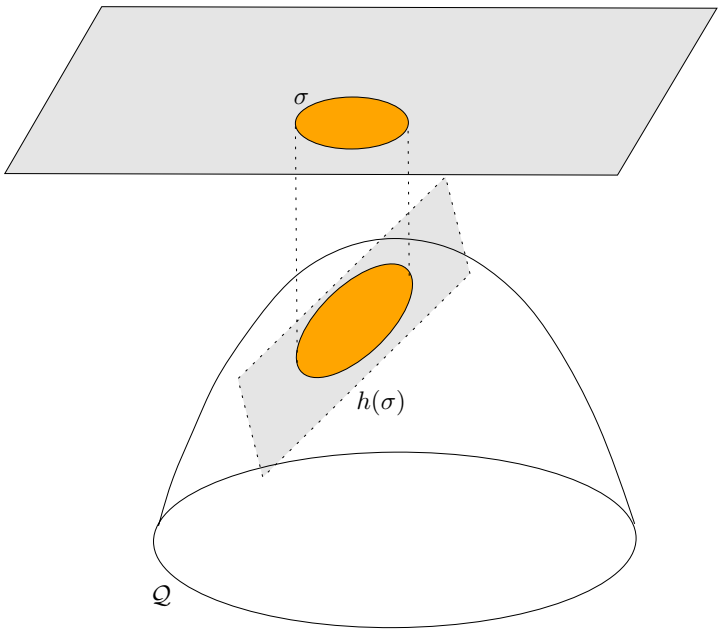

Fig. 2.5. The polar hyperplane of a sphere.

$$\phi(\sigma) \in h_{\sigma'} \iff 2c' \cdot c - s - s' = 0 \iff \phi(\sigma') \in h_\sigma$$
$$\phi(\sigma) \in h^+_{\sigma'} \iff 2c' \cdot c - s - s' > 0 \iff \phi(\sigma') \in \operatorname{int} h^+_\sigma .$$

Consider now a set $\mathcal{P} = \{p_1, \ldots, p_n\}$ of n points and let $\mathcal{V}(\mathcal{P})$ denote, as in Sect. 2.3.1, the convex polyhedron defined as the intersection of the n half-spaces below the n polar hyperplanes $h_{p_1}, \ldots, h_{p_n}$. Let f be a face of $\mathcal{V}(\mathcal{P})$ and assume that f is contained in $k+1$ hyperplanes among the h_{p_i}. Without loss of generality, we denote those hyperplanes $h_{p_1}, \ldots, h_{p_{k+1}}$. Let σ denote a hypersphere of $\mathbb{R}^d$ such that $\phi(\sigma)$ belongs to the relative interior of f. From the above discussion, we have

$$\forall i, \quad 1 \le i \le k+1, \quad \phi(\sigma) \in h_{p_i} \iff \phi(p_i) \in h_\sigma \tag{2.1}$$
$$\forall i, \ k+1 < i \le n, \quad \phi(\sigma) \in \operatorname{int} h^-_{p_i} \iff \phi(p_i) \ \in \ \operatorname{int} h^-_\sigma \tag{2.2}$$

Given a convex polyhedron $\mathcal{D}$, we say that a hyperplane h *supports* $\mathcal{D}$ if $\mathcal{D} \cap h$ is non-empty and $\mathcal{D}$ is included in one of the two halfspaces, h^+ or h^-, bounded by h. If h is a supporting hyperplane of $\mathcal{D}$, $g = \mathcal{D} \cap h$ is a face of $\mathcal{D}$. If $\mathcal{D} \subset h^-$, g is called an upper face of $\mathcal{D}$. The collection of all upper faces of $\mathcal{D}$ constitutes the *upper hull* of $\mathcal{D}$, which we denote by $\partial^+\mathcal{D}$.

Let $\mathcal{D}(\mathcal{P}) = \operatorname{conv}(\phi(\mathcal{P}))$ be the convex hull of the set $\phi(\mathcal{P})$ and consider again the face f of $\mathcal{V}(\mathcal{P})$ defined above. Write $\mathcal{P}_f = \{p_1, \ldots, p_{k+1}\}$. We deduce from (2.1) and (2.2) that, for any $\phi(\sigma)$ in the relative interior of f:

1. The hyperplane h_σ is a supporting hyperplane of $\mathcal{D}(\mathcal{P})$.
2. h_σ supports $\mathcal{D}(\mathcal{P})$ along the face $f^* = h_\sigma \cap \mathcal{D}(\mathcal{P}) = \text{conv}(\phi((\mathcal{P}_f))$.
3. $\mathcal{D}(\mathcal{P}) \subset h_\sigma^-$ and f^* is a face of $\partial^+\mathcal{D}(\mathcal{P})$.

To each face f of $\partial\mathcal{V}(\mathcal{P})$, we associate the face f^* of $\partial^+\mathcal{D}$ obtained as described above. This correspondence between the faces of $\partial\mathcal{V}(\mathcal{P})$ and the faces of $\partial^+\mathcal{D}(\mathcal{P})$ is bijective, preserves incidences and reverses inclusions, hence it is a duality.

The upper hull $\partial^+\mathcal{D}(\mathcal{P})$ projects vertically onto a cell complex of $\mathbb{R}^d$ whose vertices are the points of $\mathcal{P}$. Because the projection is 1-1, this projected cell complex is properly embedded in $\mathbb{R}^d$ and, since the projection preserves convexity, it covers the convex hull of $\mathcal{P}$. Under the general position assumption, the convex polyhedron $\mathcal{D}(\mathcal{P})$ is simplicial and the projected complex is a triangulation of $\mathcal{P}$. The duality between the faces of $\partial\mathcal{V}(\mathcal{P})$ and the faces of $\partial^+\mathcal{D}(\mathcal{P})$ implies that the projection of $\partial^+\mathcal{D}(\mathcal{P})$ is the Delaunay triangulation $\text{Del}(\mathcal{P})$ of $\mathcal{P}$ introduced at the beginning of this section. This concludes the proof that, under the general position assumption, the Delaunay triangulation $\text{Del}(\mathcal{P})$ is a triangulation of $\mathcal{P}$. We have the following diagram:

$$\begin{array}{ccc} \partial\mathcal{V}(\mathcal{P}) = \partial\left(h_{p_1}^- \cap \cdots \cap h_{p_n}^-\right) & \longleftrightarrow & \partial^+\mathcal{D}(\mathcal{P}) = \partial^+\left(\text{conv}(\phi(\mathcal{P}))\right) \\ \updownarrow & & \updownarrow \\ \text{Voronoi Diagram Vor}(\mathcal{P}) & \longleftrightarrow & \text{Delaunay Triangulation Del}(\mathcal{P}) \end{array}$$

It follows from the above correspondence that the combinatorial complexity of the Delaunay triangulation of n points is the same as the combinatorial complexity of its dual Voronoi diagram. Moreover, the Delaunay triangulation of n points of $\mathbb{R}^d$ can be deduced from the dual Voronoi diagram or vice versa in time proportional to its size. We also deduce from what precedes that computing the Delaunay triangulation of n points of $\mathbb{R}^d$ reduces to constructing the convex hull of n points of $\mathbb{R}^{d+1}$. The following theorem is then a direct consequence of known results on convex hulls [84].

Theorem 3. *The combinatorial complexity of the Voronoi diagram of n points of $\mathbb{R}^d$ and of their Delaunay triangulation is $\Theta\left(n^{\lfloor\frac{d+1}{2}\rfloor}\right)$. Both structures can be computed in optimal time $\Theta\left(n\log n + n^{\lfloor\frac{d+1}{2}\rfloor}\right)$.*

The bounds in this theorem are tight. In particular, the Voronoi diagram of n points of $\mathbb{R}^3$ may be quadratic (see Exercise 4). These bounds are worst-case bounds. Under some assumptions on the point distribution, better bounds can be obtained. For a set $\mathcal{P}$ of n points uniformly distributed in a ball of $\mathbb{R}^d$, the combinatorial complexity of the Voronoi diagram of $\mathcal{P}$ is $O(n)$ where the constant depends on the dimension d [129]. Other results are known for other point distributions [31, 33, 151].

In the discussion above, we have assumed that the points of $\mathcal{P}$ were in general position. If this is not the case, some faces of $\mathcal{D}(\mathcal{P})$ are not simplices, and the complex $\partial^+\mathcal{D}(\mathcal{P})$ projects vertically onto a cell complex, dual to the

Voronoi diagram and called the *Delaunay complex*. The faces of the *Delaunay complex* are convex and any triangulation obtained by triangulating those faces is called a Delaunay triangulation. Since there are several ways of triangulating the faces of the Delaunay complex, the Delaunay triangulation of $\mathcal{P}$ is no longer unique.

Exercise 4. Show that if we take points on two non coplanar lines of $\mathbb{R}^3$, say n_1+1 on one of the lines and n_2+1 on the other, their Voronoi diagram has $n_1 n_2$ vertices.

Exercise 5. Let S be a hypersphere of $\mathbb{R}^d$ passing through $d+1$ points $p_0, \ldots, p_d$. Show that a point p_{d+1} of $\mathbb{R}^d$ lies on S, in the interior of the ball B_S bounded by S or outside B_S, depending whether the determinant of the $(d+2)\times(d+2)$ matrix

$$\texttt{in\_sphere}(p_0,\ldots,p_{d+1}) = \begin{vmatrix} 1 & \cdots & 1 \\ p_0 & \cdots & p_{d+1} \\ p_0^2 & \cdots & p_{d+1}^2 \end{vmatrix}$$

is 0, negative or positive. This predicate is the only numerical operation that is required to check if a triangulation is a Delaunay triangulation.

Exercise 6. What are the preimages by ϕ of the points of $\mathbb{R}^{d+1}$ that lie on a line? (Distinguish the cases where the line intersects $\mathcal{Q}$ in 0, 1 or 2 points.)

Exercise 7. Project vertically the faces of the lower hull $\partial^-(\mathcal{D}(\mathcal{P}))$. Show that we obtain a triangulation of the vertices of $\mathrm{conv}(\mathcal{P})$ such that each ball circumscribing a simplex contains all the points of $\mathcal{P}$. Define a dual and make a link with Exercise 3.

Exercise 8 (Empty sphere property). Let s be any k-simplex with vertices in $\mathcal{P}$ that can be circumscribed by a a hypersphere that does not enclose any point of $\mathcal{P}$. Show that s is a face of a Delaunay triangulation of $\mathcal{P}$. Moreover, let $\mathcal{P}$ be a set of points and $\mathcal{T}$ a triangulation of $\mathcal{P}$ with the property that any hypersphere circumscribing a d-simplex of $\mathcal{T}$ does not enclose any point of $\mathcal{P}$. Show that $\mathcal{T}$ is a Delaunay triangulation of $\mathcal{P}$.

2.3.3 Power Diagrams

A construction similar to what we did for the Euclidean Voronoi diagrams of points and their dual Delaunay triangulations can be done for the so-called power or Laguerre diagrams. Here we take as our finite set of objects a set of hyperspheres (instead of points) and consider as distance function of a point x to a hypersphere σ the power of x to σ. As we will see, the class of power diagrams is identical to the class of affine diagrams, i.e. the diagrams whose bisectors are hyperplanes.

Definition of Power Diagrams

We call *power* of a point x to a hypersphere σ of center c and radius r the real number

$$\sigma(x) = (x - c)^2 - r^2.$$

Let $\mathcal{S} = \{\sigma_1, \ldots, \sigma_n\}$ be a set of hyperspheres of $\mathbb{R}^d$. We denote by c_i the center of σ_i, r_i its radius, $\sigma_i(x) = (x - c_i)^2 - r_i^2$ the power function to σ_i, and $s_i = c_i^2 - r_i^2$ the power of the origin. To each σ_i, we associate the region $L(\sigma_i)$ consisting of the points of $\mathbb{R}^d$ whose power to σ_i is not larger than their power to the other hyperspheres of $\mathcal{S}$:

$$L(\sigma_i) = \{x \in \mathbb{R}^d : \sigma_i(x) \leq \sigma_j(x), 1 \leq j \leq n\}.$$

The set of points that have equal power to two hyperspheres σ_i and σ_j is a hyperplane, noted π_{ij}, called *the radical hyperplane* of σ_i and σ_j. Hyperplane π_{ij} is orthogonal to the line joining the centers of σ_i and σ_j. We denote by π_{ij}^i the half-space bounded by π_{ij} consisting of the points whose power to σ_i is smaller than their power to σ_j. The region $L(\sigma_i)$ is the intersection of all half-spaces π_{ij}^i, $j \neq i$. If this intersection is not empty, it is a convex polyhedron, possibly not bounded. We call *power regions* the non empty regions $L(\sigma_i)$.

We define the *power diagram* of $\mathcal{S}$, noted $\mathrm{Pow}(\mathcal{S})$, as the cell complex whose cells are the power regions and their faces. When all hyperspheres have the same radius, their power diagram is identical to the Voronoi diagram of their centers.

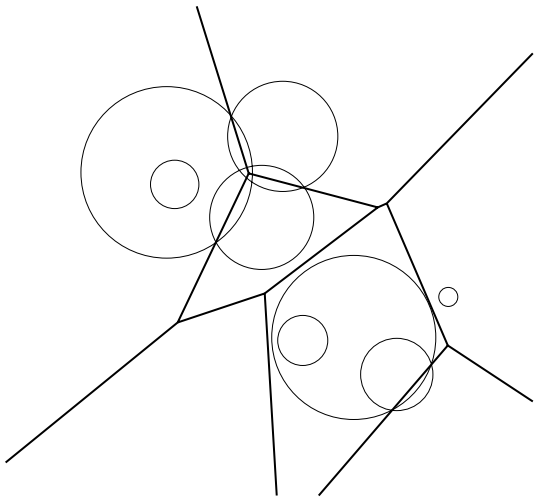

Fig. 2.6. A power diagram

Equivalently, the power diagram of $\mathcal{S}$ can be defined as the minimization diagram of the functions $\sigma_i, \ldots, \sigma_n$. Observing that for any x

$$\arg\min_i \sigma_i(x) = \arg\min_i(-2c_i \cdot x + s_i),$$

we obtain that the power diagram of $\mathcal{S}$ is the minimization diagram of the set of affine functions

$$d_i(x) = -2p_i \cdot x + s_i$$

whose graphs are hyperplanes of $\mathbb{R}^{d+1}$. Let us call h_{σ_i}, $i = 1, \ldots, n$, those hyperplanes and let $h^-_{\sigma_i}$ denote the half-space lying below h_{σ_i}. The minimization diagram of the δ_i is obtained by projecting vertically the convex polyhedron

$$\mathcal{L}(\mathcal{S}) = h^-_{p_1} \cap \cdots \cap h^-_{p_n}.$$

Theorem 4. *The faces of the power diagram* $\mathrm{Pow}(\mathcal{S})$ *of* $\mathcal{S}$ *are the vertical projections of the faces of the convex polyhedron* $\mathcal{L}(\mathcal{S})$.

Power diagrams are very similar to Voronoi diagrams: the only difference is that the hyperplanes supporting the faces of $\mathcal{L}(\mathcal{S})$ are not necessarily tangent to the paraboloid $\mathcal{Q}$ and that some hyperplane may not contribute a face. In other words, some hypersphere σ_i may have an empty power region (see the small circle in the upper left corner of Fig. 2.6).

By proceeding as in Sect. 2.3.2, we can define a convex polyhedron $\mathcal{R}(\mathcal{S})$ whose upper hull $\partial^+\mathcal{R}(\mathcal{S})$ is dual to $\partial\mathcal{L}(\mathcal{S})$. The vertical projection of the faces of $\partial^+\mathcal{R}(\mathcal{S})$ constitute the faces of a cell complex which, in general, is a simplicial complex. We call such a complex the *regular triangulation* of $\mathcal{S}$ and denote it by $\mathrm{Reg}(\mathcal{S})$. We have the following diagram:

$$\begin{array}{ccc} \partial\mathcal{L}(\mathcal{S}) = \partial\left(h^-_{\sigma_1} \cap \cdots \cap h^-_{\sigma_n}\right) & \longleftrightarrow & \partial^+\mathcal{R}(\mathcal{S}) = \partial^+\mathrm{conv}(\phi(\mathcal{S})) \\ \updownarrow & & \updownarrow \\ \text{Power diagram } \mathrm{Pow}(\mathcal{S}) & \longleftrightarrow & \text{Regular triangulation } \mathrm{Reg}(\mathcal{S}) \end{array}$$

We deduce the following theorem that states that computing the power diagram of n hyperspheres of $\mathbb{R}^d$ (or equivalently its dual regular triangulation) has the same asymptotic complexity as computing the Euclidean Voronoi diagram or the Delaunay triangulation of n points of $\mathbb{R}^d$.

Theorem 5. *The combinatorial complexity of the power diagram of* n *hyperspheres of* $\mathbb{R}^d$ *and of its dual regular triangulation are* $\Theta\left(n^{\lfloor\frac{d+1}{2}\rfloor}\right)$. *Both structures can be computed in optimal time* $\Theta\left(n\log n + n^{\lfloor\frac{d+1}{2}\rfloor}\right)$.

Affine Voronoi Diagrams

Euclidean Voronoi diagrams of points and power diagrams of hyperspheres are two examples of minimization diagrams whose bisectors are hyperplanes. It is interesting to classify Voronoi diagrams with respect to their bisectors. A first important class of Voronoi diagrams is the class of *affine diagrams* which consists of all Voronoi diagrams whose bisectors are hyperplanes.

In Sect. 2.5, we will prove that any affine Voronoi diagram of $\mathbb{R}^d$ is identical to the power diagram of some set of hyperspheres of $\mathbb{R}^d$ (Theorem 13), therefore showing that the class of affine Voronoi diagrams is identical to the class of power diagrams.

Exercise 9. Show that the intersection of a power diagram with an affine subspace is still a power diagram and compute the corresponding spheres.

Exercise 10. Show that any power diagram of $\mathbb{R}^d$ is the intersection of a Voronoi diagram of $\mathbb{R}^{d+1}$ by a hyperplane.

Exercise 11. Show that the only numerical operation that is required to check if a triangulation is the regular triangulation of a set of hyperspheres σ_i is the evaluation of the sign of the determinant of the $(d+2) \times (d+2)$ matrix

$$\texttt{power\_test}(\sigma_0, \ldots, \sigma_{d+1}) = \begin{vmatrix} 1 & \cdots & 1 \\ c_0 & \cdots & c_{d+1} \\ c_0^2 - r_0^2 & \cdots & c_{d+1}^2 - r_{d+1}^2 \end{vmatrix}$$

where c_i and r_i are respectively the center and the radius of σ_i.

2.4 Voronoi Diagrams with Algebraic Bisectors

In this section, we introduce a first class of non-affine diagrams, namely the class of diagrams whose bisectors are algebraic hypersurfaces. We first consider the case of Möbius diagrams whose bisectors are hyperspheres and the case of *anisotropic* diagrams whose bisectors are quadratic hypersurfaces. These diagrams can be computed through linearization, a technique to be described in full generality in Sect. 2.5. Apollonius (or Johnson-Mehl) diagrams, although semi-algebraic and not algebraic, are also described in this section since they are closely related to Möbius diagrams and can also be linearized.

2.4.1 Möbius Diagrams

In this section, we introduce a class of non-affine Voronoi diagrams, the so-called Möbius diagrams, introduced by Boissonnat and Karavelas [63].

The class of Möbius diagrams includes affine diagrams. In fact, as we will see, the class of Möbius diagrams is identical to the class of diagrams whose bisectors are hyperspheres (or hyperplanes).

Definition of Möbius Diagrams

Let $\omega = \{\omega_1, \ldots, \omega_n\}$ be a set of so-called *Möbius sites* of $\mathbb{R}^d$, where ω_i is a triple (p_i, λ_i, μ_i) formed of a point p_i of $\mathbb{R}^d$, and two real numbers λ_i and μ_i.

For a point $x \in \mathbb{R}^d$, the distance $\delta_i(x)$ from x to the Möbius site ω_i is defined as

$$\delta_i(x) = \lambda_i(x - p_i)^2 - \mu_i.$$

Observe that the graph of δ_i is a paraboloid of revolution whose axis is vertical. The Möbius region of the Möbius site ω_i, $i = 1, \ldots, n$, is

$$M(\omega_i) = \{x \in \mathbb{R}^d : \delta_i(x) \leq \delta_j(x), 1 \leq j \leq n\}.$$

Observe that a Möbius region may be non-contractible and even disconnected.

The minimization diagram of the δ_i is called the *Möbius diagram* of ω and noted Möb (ω) (see Fig. 2.4.1).

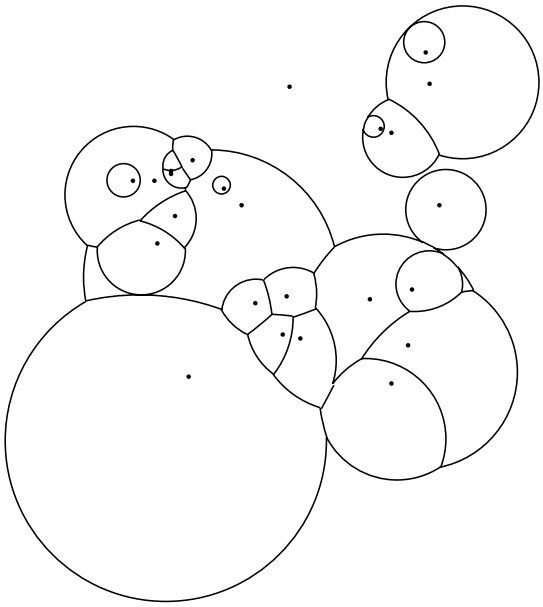

Fig. 2.7. A Möbius diagram

Möbius diagrams are generalizations of Euclidean Voronoi and power diagrams. In particular, if all λ_i are equal to some positive λ, the Möbius diagram coincides with the power diagram of a set of spheres $\{\sigma_i, i = 1, \ldots, n\}$, where σ_i is the sphere centered at p_i of squared radius μ_i/λ. If all μ_i are equal and all λ_i are positive, then the Möbius diagram coincides with the so-called *multiplicatively weighted Voronoi diagram* of the weighted points $(p_i, \sqrt{\lambda_i})$.

The following lemma states that the bisector of two Möbius sites is a hypersphere (possibly degenerated in a point or in a hyperplane). Its proof is straightforward.

Lemma 1. *Let $\omega_i = \{p_i, \lambda_i, \mu_i\}$ and $\omega_j = \{p_j, \lambda_j, \mu_j\}$, $\omega_i \neq \omega_j$ be two Möbius sites. The bisector σ_{ij} of ω_i and ω_j is the empty set, a single point, a hypersphere or a hyperplane.*

Möbius Diagrams and Power Diagrams

We now present an equivalence between Möbius diagrams in $\mathbb{R}^d$ and power diagrams in $\mathbb{R}^{d+1}$. This result is a direct generalization of a similar result for multiplicatively weighted diagrams [35]. Given a cell complex $\mathcal{C}$ covering a subspace X, we call *restriction of $\mathcal{C}$ to X* the subdivision of X whose faces are the intersections of the faces of $\mathcal{C}$ with X. The restriction of $\mathcal{C}$ to X is denoted by $\mathcal{C}_X$. Note that the restriction $\mathcal{C}_X$ is not, in general, a cell complex and that its faces may be non-contractible and even non-connected.

We associate to $\boldsymbol{\omega}= \{\omega_1, \ldots, \omega_n\}$ the set of hyperspheres $\Sigma = \{\Sigma_1, \ldots, \Sigma_n\}$ of $\mathbb{R}^{d+1}$ of equations

$$\Sigma_i(X) = X^2 - 2C_i \cdot X + s_i = 0,$$

where $C_i = (\lambda_i p_i, -\frac{\lambda_i}{2})$ and $s_i = \lambda_i\, p_i^2 - \mu_i$. We denote by $\mathcal{Q}$ the paraboloid of $\mathbb{R}^{d+1}$ of equation $x_{d+1} - x^2 = 0$.

Theorem 6 (Linearization). *The Möbius diagram* $\mathrm{Möb}(\boldsymbol{\omega})$ *of $\boldsymbol{\omega}$ is obtained by projecting vertically the faces of the restriction* $\mathrm{Pow}_{\mathcal{Q}}(\Sigma)$ *of the power diagram of Σ to $\mathcal{Q}$.*

Proof. If $x \in \mathbb{R}^d$ is closer to ω_i than to ω_j with respect to δ_M, we have for all $j = 1, \ldots, n$,

$$\begin{aligned}
&\lambda_i(x - p_i)^2 - \mu_i \leq \lambda_j(x - p_j)^2 - \mu_j \\
&\quad\iff \lambda_i x^2 - 2\lambda_i p_i \cdot x + \lambda_i p_i^2 - \mu_i \leq \lambda_j x^2 - 2\lambda_j p_j \cdot x + \lambda_j p_j^2 - \mu_j \\
&\quad\iff (x^2 + \tfrac{\lambda_i}{2})^2 + (x - \lambda_i p_i)^2 - \tfrac{\lambda_i^2}{4} - \lambda_i^2 p_i^2 + \lambda_i p_i^2 - \mu_i \\
&\qquad\leq (x^2 + \tfrac{\lambda_j}{2})^2 + (x - \lambda_j p_j)^2 - \tfrac{\lambda_j^2}{4} - \lambda_j^2 p_j^2 + \lambda_j p_j^2 - \mu_j \\
&\quad\iff (X - C_i)^2 - r_i^2 \leq (X - C_j)^2 - r_j^2 \\
&\quad\iff \Sigma_i(X) \leq \Sigma_j(X)
\end{aligned}$$

where $X = (x, x^2) \in \mathcal{Q} \subset \mathbb{R}^{d+1}$, $C_i = (\lambda_i p_i, -\frac{\lambda_i}{2}) \in \mathbb{R}^{d+1}$ and $r_i^2 = \lambda_i^2 p_i^2 + \frac{\lambda_i^2}{4} - \lambda_i p_i^2 + \mu_i$. The above inequality shows that x is closer to ω_i than to ω_j if and only if X belongs to the power region of Σ_i in the power diagram of the hyperspheres Σ_j, $j = 1, \ldots, n$. As X belongs to $\mathcal{Q}$ and projects vertically onto x, we have proved the result.

Corollary 1. *Let Σ be a finite set of hyperspheres of $\mathbb{R}^{d+1}$,* $\mathrm{Pow}(\Sigma)$ *its power diagram and* $\mathrm{Pow}_{\mathcal{Q}}(\Sigma)$ *the restriction of* $\mathrm{Pow}(\Sigma)$ *to $\mathcal{Q}$. The vertical projection of* $\mathrm{Pow}_{\mathcal{Q}}(\Sigma)$ *is the Möbius diagram* $\mathrm{Möb}(\boldsymbol{\omega})$ *of a set of Möbius sites of $\mathbb{R}^d$.*

Easy computations give $\boldsymbol{\omega}$.

Combinatorial and Algorithmic Properties

It follows from Theorem 6 that the combinatorial complexity of the Möbius diagram of n Möbius sites in $\mathbb{R}^d$ is $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$. This bound is tight since Aurenhammer [35] has shown that it is tight for multiplicatively weighted Voronoi diagrams.

We easily deduce from the proof of the Linearization Theorem 6 an algorithm for constructing Möbius diagrams. First, we compute the power diagram of the hyperspheres Σ_i of $\mathbb{R}^{d+1}$, intersect each of the faces of this diagram with the paraboloid $\mathcal{Q}$ and then project the result on $\mathbb{R}^d$.

Theorem 7. *Let ω be a set of n Möbius sites in $\mathbb{R}^d$, $d \geq 2$. The Möbius diagram* $\mathrm{Möb}(\omega)$ *of ω can be constructed in worst-case optimal time $\Theta(n \log n + n^{\lfloor \frac{d}{2} \rfloor + 1})$.*

Another consequence of the linearization theorem is the fact that any Möbius diagram can be represented as a simplicial complex $T_{\mathcal{Q}}$ embedded in $\mathbb{R}^{d+1}$. $T_{\mathcal{Q}}$ is a sub-complex of the regular triangulation T dual to the power diagram $\mathrm{Pow}(\Sigma_i)$ of the hyperspheres Σ_i. Since T is embedded in $\mathbb{R}^{d+1}$, $T_{\mathcal{Q}}$ is a simplicial complex of $\mathbb{R}^{d+1}$. More precisely, $T_{\mathcal{Q}}$ consists of the faces of T that are dual to the faces of $\mathrm{Pow}_{\mathcal{Q}}(\Sigma)$, i.e. the faces of the power diagram that intersect $\mathcal{Q}$. We will call $T_{\mathcal{Q}}$ the *dual* of $\mathrm{Pow}_{\mathcal{Q}}(\Sigma)$. Observe that since, in general, no vertex of $\mathrm{Pow}(\Sigma)$ lies on $\mathcal{Q}$, $T_{\mathcal{Q}}$ is a d-dimensional simplicial complex (embedded in $\mathbb{R}^{d+1}$).

Moreover, if the faces of $\mathrm{Pow}(\Sigma)$ intersect $\mathcal{Q}$ transversally and along topological balls, then, by a result of Edelsbrunner and Shah [138], $T_{\mathcal{Q}}$ is homeomorphic to $\mathcal{Q}$ and therefore to $\mathbb{R}^d$. It should be noted that this result states that the simplicial complex $T_{\mathcal{Q}}$ has the topology of $\mathbb{R}^d$. This result, however, is mainly combinatorial, and does not imply that the embedding of $T_{\mathcal{Q}}$ into $\mathbb{R}^{d+1}$ as a sub-complex of the regular triangulation T may be projected in a 1-1 manner onto $\mathbb{R}^d$.

Spherical Voronoi Diagrams

Lemma 1 states that the bisectors of two Möbius sites is a hypersphere (possibly degenerated in a hyperplane). More generally, let us consider the Voronoi diagrams such that, for any two objects o_i and o_j of $\mathcal{O}$, the bisector $\sigma_{ij} = \{x \in \mathbb{R}^d, \delta_i(x) = \delta_j(x)\}$ is a hypersphere. Such a diagram is called a *spherical Voronoi diagram.*

In Sect. 2.5, we will prove that any spherical Voronoi diagram of $\mathbb{R}^d$ is a Möbius diagram (Theorem 15).

Möbius transformations are the transformations that preserve hyperspheres. An example of a Möbius transformation is the inversion with respect to a hypersphere. If the hypersphere is centered at c and has radius r, the inversion associates to a point $x \in \mathbb{R}^d$ its image

$$x' = c + \frac{r(x-c)}{(x-c)^2}.$$

Moreover, it is known that any Möbius transformation is the composition of up to four inversions [108]. An immediate consequence of Theorem 15 is that the set of Möbius diagrams in $\mathbb{R}^d$ is stable under Möbius transformations, hence their name.

Möbius Diagrams on Spheres

Given a set ω of n Möbius sites of $\mathbb{R}^{d+1}$, the restriction of their Möbius diagram to a hypersphere $\mathbb{S}^d$ is called a Möbius diagram on $\mathbb{S}^d$. It can be shown that such a diagram is also the restriction of a power diagram of hyperspheres of $\mathbb{R}^{d+1}$ to $\mathbb{S}^d$ (Exercise 14) .

We define spherical diagrams on $\mathbb{S}^d$ as the diagrams on $\mathbb{S}^d$ whose bisectors are hyperspheres of $\mathbb{S}^d$ and that satisfy two properties detailed in Sect. 2.5.1 and 2.5.2. See Exercise 16 for more details on these conditions. This exercise proves that the restriction of a Möbius diagram, i.e. a Möbius diagram on $\mathbb{S}^d$, is a spherical diagram.

Let us now prove the converse: any spherical diagram on $\mathbb{S}^d$ is a Möbius diagram on $\mathbb{S}^d$. Let h be a hyperplane of $\mathbb{R}^{d+1}$. The stereographic projection that maps $\mathbb{S}^d$ to h maps any spherical diagram $\mathcal{D}$ on $\mathbb{S}^d$ to some spherical diagram $\mathcal{D}'$ on h. Theorem 15 implies that this $\mathcal{D}'$ is in fact a Möbius diagram. Exercise 12 shows that $\mathcal{D}$, which is the image of $\mathcal{D}'$ by the inverse of the stereographic projection, is the restriction of some power diagram of $\mathbb{R}^{d+1}$ to $\mathbb{S}^d$. Exercise 14 then proves that it is indeed a Möbius diagram.

Exercise 12. Show that the linearization theorem and its corollary still hold if one replaces the paraboloid $\mathcal{Q}$ by any hypersphere of $\mathbb{R}^{d+1}$ and the vertical projection by the corresponding stereographic projection.

Exercise 13. Show that the intersection of a Möbius diagram in $\mathbb{R}^d$ with a k-flat or a k-sphere σ is a Möbius diagram in σ.

Exercise 14. Show that the restriction of a Möbius diagram of n Möbius sites to a hypersphere $\Sigma \subset \mathbb{R}^d$ (i.e. a Möbius diagram on Σ) is identical to the restriction of a power diagram of n hyperspheres of $\mathbb{R}^d$ with Σ, and vice versa.

Exercise 15. The predicates needed for constructing a Möbius diagram are those needed to construct $\mathrm{Pow}(\Sigma)$ and those that decide whether a face of $\mathrm{Pow}(\Sigma)$ intersects $\mathcal{Q}$ or not. Write the corresponding algebraic expressions.

Exercise 16. Explain how the two conditions A.C. and L.C.C. presented in Sect. 2.5.1 and 2.5.2 are to be adapted to the case of spherical diagrams on a sphere (Hint: consider the L.C.C. condition as a pencil condition, and define

a pencil of circles on a sphere as the intersection of a pencil of hyperplanes with this sphere).

Note that the restriction of a Voronoi diagram (affine or not) to $\mathbb{S}^d$ always satisfies this adapted version of A.C. and prove that the restriction of an affine Voronoi diagram to $\mathbb{S}^d$ satisfies L.C.C. so that Exercise 14 allows to conclude that the restriction of a Möbius diagram to $\mathbb{S}^d$, i.e. Möbius diagram on $\mathbb{S}^d$, is a spherical diagram.

2.4.2 Anisotropic Diagrams

The definition of anisotropic Voronoi diagrams presented in this section is a slight extension of the definition proposed by Labelle and Shewchuk [234]. The objects are points and the distance to a point is a quadratic form with an additive weight.

Anisotropic diagrams appear to be a natural generalization of Möbius diagrams and reduce to Möbius diagrams when the matrices are taken to be a scalar times the identity matrix. As will be shown, the class of anisotropic diagrams is identical to the class of diagrams whose bisectors are quadratic hypersurfaces.

Definition and linearization

Consider a finite set of anisotropic sites $S = \{s_1, \ldots, s_n\}$. Each site s_i, $i = 1, \ldots, n$, is a triple (p_i, M_i, π_i) formed by a point $p_i \in \mathbb{R}^d$, a $d \times d$ symmetric positive definite matrix M_i and a scalar weight π_i. The distance $\delta_i(x)$ of point $x \in \mathbb{R}^d$ to site s_i is defined by

$$\delta_i(x) = (x - p_i)^t M_i (x - p_i) - \pi_i.$$

The anisotropic Voronoi region of site s is then defined as

$$AV(s_i) = \{x \in \mathbb{R}^d, \delta_i(x) \leq \delta_j(x), \forall 1 \leq j \leq n\},$$

The *anisotropic Voronoi diagram* is the minimization diagram of the functions $\delta_i(x)$.

Let $D = \frac{d(d+3)}{2}$. To each point $x = (x_1, \ldots, x_d) \in \mathbb{R}^d$, we associate the two points

$$\tilde{\phi}(x) = (x_r x_i, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}}$$
$$\hat{\phi}(x) = (x, \tilde{\phi}(x)) \in \mathbb{R}^D,$$

and we denote by $\mathcal{Q}$ the d-manifold of $\mathbb{R}^D$ defined as

$$\mathcal{Q} = \left\{ \hat{\phi}(x),\ x \in \mathbb{R}^d \right\}.$$

To each site $s_i = (p_i, M_i, \pi_i) \in S$, we associate:

1. the point $\tilde{m}_i \in \mathbb{R}^{\frac{d(d+1)}{2}}$ defined as

$$\tilde{m}_i^{u,u} = -\frac{1}{2} M_i^{u,u}, \text{ for } 1 \leq u \leq d;$$
$$\tilde{m}_i^{u,v} = -M_i^{u,v}, \text{ for } 1 \leq u < v \leq d,$$

2. the point $\hat{p}_i = (M_i p_i, \tilde{m}_i)$,
3. the sphere σ_i of center $\hat{p}_i$ and radius $\sqrt{\|\hat{p}_i\|^2 - p_i^t M_i p_i - \pi_i}$.

Let Π be the projection $\hat{y} = (y, \tilde{y}) \in \mathbb{R}^D \mapsto y \in \mathbb{R}^d$ and let Σ be the set of spheres σ_i, $i = 1, \ldots, n$.

Theorem 8 (Linearization). *The anisotropic diagram of S is the image by Π of the restriction of the power diagram* $\mathrm{Pow}(\Sigma)$ *to the d-manifold $\mathcal{Q}$.*

Proof. We have

$$\begin{aligned}\delta_i(x) &= (x - p_i)^t M (x - p_i) - \pi_i \\ &= x^t M_i x - 2 p_i^t M_i x + p_i^t M_i p_i - \pi_i \\ &= -2 \hat{p}_i^t \hat{\phi}(x) + p_i^t M_i p_i - \pi_i\end{aligned}$$

This implies that $\delta_i(x) < \delta_j(x)$ if and only if

$$(\hat{\phi}(x) - \hat{p}_i)^2 - (\hat{p}_i^2 - p_i^t M_i p_i - \pi_i) < (\hat{\phi}(x) - \hat{p}_j)^2 - (\hat{p}_j^2 - p_j^t M_j p_j - \pi_j)$$

Hence, x is closer to s_i than to s_j if and only if the power of $\hat{\phi}(x)$ to σ_i is smaller than its power to σ_j. Equivalently, a point $\hat{\phi}(x) \in \mathcal{Q}$ belongs to the power cell of $\sigma(s_i)$ if and only if its projection $x = \Pi(\hat{\phi}(x))$ belongs to the anisotropic Voronoi region $AV(s_i)$.

We easily deduce the following theorem.

Theorem 9. *The Voronoi diagram of n anisotropic sites of $\mathbb{R}^d$ can be computed in time $O(n^{\lfloor \frac{D+1}{2} \rfloor})$ where $D = \frac{d(d+3)}{2}$.*

This result is to be compared to Theorem 1 which provides a better combinatorial bound. We let as an open question to fill the gap between those two bounds.

Quadratic Voronoi Diagrams

The bisectors of anisotropic diagrams, as defined in the previous section, are quadratic hypersurfaces. A minimization diagram whose bisectors are hyperquadrics is called a *quadratic Voronoi diagram*. In Sect. 2.5, we will prove that any quadratic Voronoi diagram is the anisotropic Voronoi diagram of a set of anisotropic sites (Theorem 16).

2.4.3 Apollonius Diagrams

In this section, we present diagrams that are closely related to Möbius diagrams: namely, the Euclidean Voronoi diagrams of hyperspheres, also called Apollonius or Johnson-Mehl diagrams. Contrary to Möbius and anisotropic diagrams, the bisectors of Apollonius diagrams are not algebraic hypersurfaces since the bisector between two hyperspheres is only one sheet of a hyperboloid. As a consequence, Apollonius diagrams cannot be linearized in the same way as Möbius and anisotropic diagrams. Nevertheless, another linearization scheme can be applied, leading to interesting combinatorial and algorithmic results.

Definition of Apollonius Diagrams

Let us consider a finite set of weighted points $\mathcal{S} = \{\sigma_0, \sigma_1, \ldots, \sigma_n\}$ where $\sigma_i = (p_i, r_i)$, $p_i \in \mathbb{R}^d$ and $r_i \in \mathbb{R}$. We define the distance from x to σ_i as

$$\delta_i(x) = \|x - p_i\| - r_i.$$

This distance is also called the *additively weighted* distance from x to the weighted point σ_i. The minimization diagram of the distance functions δ_i, $i = 1, \ldots, n$, is called the additively weighted Voronoi diagram, or the Apollonius diagram of $\mathcal{S}$. We denote it by $\mathrm{Apo}(\mathcal{S})$ (see Fig. 2.8).

The Apollonius region $A(\sigma_i)$ of σ_i is defined as

$$A(\sigma_i) = \{x \in \mathbb{R}^d, \delta_i(x) \leq \delta_j(x)\}.$$

It is easy to see that $A(\sigma_i)$ is either empty or star-shaped from p_i. The boundary of $A(\sigma_i)$ may have a complicated structure. In fact, as we will see, the boundary of $A(\sigma_i)$ has the same combinatorial structure as a Möbius diagram in $\mathbb{R}^{d-1}$.

Since the diagram is not changed if we replace all r_i by $r_i + r$ for any $r \in \mathbb{R}$, we can assume, without loss of generality, that all r_i are non negative. The weighted points are then hyperspheres and the distance to a weighted point is the signed Euclidean distance to the corresponding hypersphere, counted positively outside the hypersphere and negatively inside the hypersphere.

Observe that, in the plane, a vertex of an Apollonius diagram is the center of a circle tangent to three circles of $\mathcal{S}$ (assuming all r_i non negative). Computing such a point is known as Apollonius' Tenth Problem, hence the name of the diagram.

Apollonius Diagrams and Power Diagrams

The graph of the distance function $\delta_i(x)$ is the half-cone of revolution $\mathcal{C}_i$ of equation

$$\mathcal{C}_i : x_{d+1} = \|x - p_i\| - r_i, \quad x_{d+1} + r_i \geq 0$$

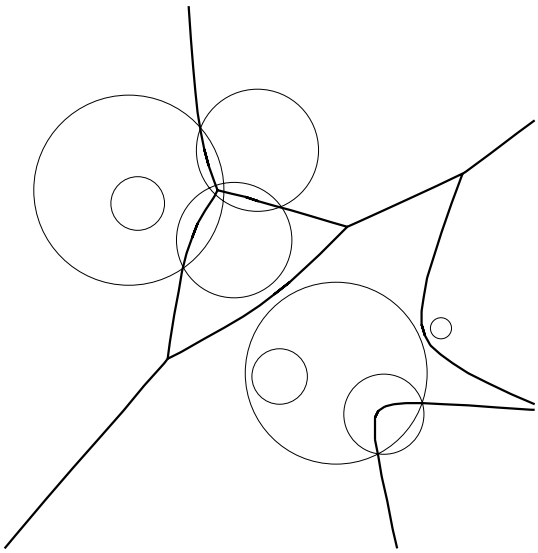

Fig. 2.8. The Apollonius diagram of a set of circles. Compare with the power diagram of the same set of circles in Fig. 2.6

The bisector of two hyperspheres of $\mathcal{S}$ is thus the projection of the intersection of two half-cones. This intersection is a quadratic hypersurface (in fact, a sheet of a two sheet hyperboloid) contained in a hyperplane. Indeed, we have

$$\begin{aligned}\mathcal{C}_1 \ &: \ (x_{d+1}+r_1)^2 = (x-p_1)^2, \ \ x_{d+1}+r_1 > 0,\\ \mathcal{C}_2 \ &: \ (x_{d+1}+r_2)^2 = (x-p_2)^2, \ \ x_{d+1}+r_2 > 0.\end{aligned}$$

The intersection of the two half-cones is contained in the hyperplane h_{12} whose equation is obtained by subtracting the two sides of the above equations:

$$h_{12} \ \ : \ \ -2(p_1-p_2)\cdot x - 2(r_1-r_2)x_{d+1} + p_1^2 - r_1^2 - p_2^2 + r_2^2 = 0.$$

This shows that there exists a correspondence between the diagram $\mathrm{Apo}(\mathcal{S})$ and the power diagram of the hyperspheres Σ_i in $\mathbb{R}^{d+1}$ $(i = 1, \ldots, n)$, where Σ_i is centered at (p_i, r_i) and has radius $r_i\sqrt{2}$. More precisely, $A(\sigma_i)$ is the projection of the intersection of the half-cone $\mathcal{C}_i$ with the power region $L(\Sigma_i)$. Indeed, x is in $A(\sigma_i)$ if and only if the projection X_i of x onto $\mathcal{C}_i$ has a smaller x_{d+1}-coordinate than the projections of x onto the other half-cones $\mathcal{C}_j$, $j \neq i$. In other words, the coordinates (x, x_{d+1}) of X_i must obey

$$\begin{aligned}&(x_{d+1}+r_i)^2 = (x-p_i)^2\\ &(x_{d+1}+r_j)^2 \leq (x-p_j)^2 \quad \text{for any} \ \ j \neq i,\end{aligned}$$

and by subtracting both sides, it follows that $\Sigma_i(X_i) \leq \Sigma_j(X_i)$ for all j.

Algorithm 2 Construction of Apollonius diagrams

INPUT: a set of hyperspheres $\mathcal{S}$

1. Compute Σ_i, for $i = 1, \ldots, n$;
2. Compute the power diagram of the Σ_i's;
3. For all $i = 1, \ldots, n$, project vertically the intersection of the power region $L(\Sigma_i)$ with the half-cone $\mathcal{C}_i$.

OUTPUT: the Apollonius diagram of $\mathcal{S}$.

The Apollonius diagram of $\mathcal{S}$ can be computed using the following algorithm:

The power diagram of the Σ_i can be computed in time $O(n^{\lfloor \frac{d}{2} \rfloor + 1} \log n)$. The intersection involved in Step 3 can be computed in time proportional to the number of faces of the power diagram of the Σ_i's, which is $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$. We have thus proved the following theorem due to Aurenhammer [35]:

Theorem 10. *The Apollonius diagram of a set of n hyperspheres in $\mathbb{R}^d$ has complexity $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$ and can be computed in time $O(n^{\lfloor \frac{d}{2} \rfloor + 1} \log n)$.*

This result is optimal in odd dimensions, since the bounds above coincide with the corresponding bounds for the Voronoi diagram of points under the Euclidean distance. It is not optimal in dimension 2 (see Exercise 20). We also conjecture that it is not optimal in any even dimension.

Computing a Single Apollonius Region

We now establish a correspondence, due to Boissonnat and Karavelas [63], between a single Apollonius region and a Möbius diagram on a hypersphere.

To give the intuition behind the result, we consider first the case where one of the hyperspheres, say σ_0, is a hyperplane, i.e. a hypersphere of infinite radius. We take for σ_0 the hyperplane $x_d = 0$, and assume that all the other hyperspheres lie the half-space $x_d > 0$. The distance $\delta_0(x)$ from a point $x \in \mathbb{R}^d$ to σ_0 is defined as the Euclidean distance.

The points that are at equal distance from σ_0 and σ_i, $i > 0$, belong to a paraboloid of revolution with vertical axis. Consider such a paraboloid as the graph of a $(d-1)$-variate function ϑ_i defined over $\mathbb{R}^{d-1}$. If follows from Sect. 2.4.1 that the minimization diagram of the ϑ_i, $i = 1, \ldots, n$, is a Möbius diagram (see Fig. 2.9).

Easy computations give the associated weighted points. Write $p_i = (p_i', p_i'')$, $p_i' \in \mathbb{R}^{d-1}$, $p_i'' \in \mathbb{R}$, $i > 0$ and let $\boldsymbol{\omega} = \{\omega_1, \ldots, \omega_n\}$ be the set of Möbius sites of $\mathbb{R}^d$ where $\omega_i = \{p_i', \lambda_i, \mu_i\}$, and

$$\lambda_i = \frac{1}{r_i + p_i''}, \qquad \mu_i = r_i - p_i'', \qquad i > 0.$$

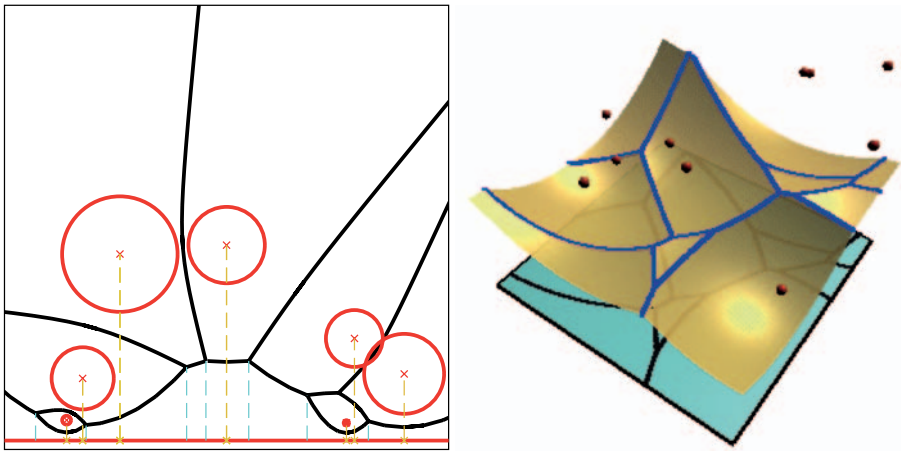

Fig. 2.9. A cell in an Apollonius diagram of hyperspheres projects vertically onto a Möbius diagram in σ_0

We let as an exercise to verify that the vertical projection of the boundary of the Apollonius region $A(\sigma_0)$ of σ_0 onto σ_0 is the Möbius diagram of ω.

We have assumed that one of the hyperspheres was a hyperplane. We now consider the case of hyperspheres of finite radii. The crucial observation is that the radial projection of $A(\sigma_0) \cap A(\sigma_i) \cap A(\sigma_j)$ onto σ_0, if not empty, is a hypersphere. It follows that the radial projection of the boundary of $A(\sigma_0)$ onto σ_0 is a Möbius diagram on σ_0.

Such a Möbius diagram on σ_0 can be computed by constructing the restriction of the power diagram of n hyperspheres of $\mathbb{R}^d$ with the hypersphere σ_0 (see Exercise 14).

Theorem 11. *Let $\mathcal{S}$ be a set of n hyperspheres in $\mathbb{R}^d$. The worst-case complexity of a single Apollonius region in the diagram of n hyperspheres of $\mathbb{R}^d$ is $\Theta(n^{\lfloor \frac{d+1}{2} \rfloor})$. Such a region can be computed in optimal time $\Theta(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor})$.*

Exercise 17. Show that the cell of hypersphere σ_i in the Apollonius diagram of $\mathcal{S}$ is empty if and only if σ_i is inside another hypersphere σ_j.

Exercise 18. The predicates required to construct an Apollonius region are multivariate polynomials of degree at most 8 and 16 when $d = 2$ and 3 respectively. Detail these predicates [62].

Exercise 19. Show that the convex hull of a finite number of hyperspheres can be deduced from the restriction of a power diagram to a unit hypersphere [62].

Exercise 20. Prove that the combinatorial complexity of the Apollonius diagram of n circles in the plane has linear size.

Exercise 21 (Open problem). Give a tight bound on the combinatorial complexity of the Apollonius diagram of n hyperspheres of $\mathbb{R}^d$ when d is even.

2.5 Linearization

In this section, we introduce abstract diagrams, which are diagrams defined in terms of their bisectors. We impose suitable conditions on these bisectors so that any abstract diagram can be built as the minimization diagram of some distance functions, thus showing that the class of abstract diagrams is the same as the class of Voronoi diagrams.

Given a class of bisectors, such as affine or spherical bisectors, we then consider the inverse problem of determining a small class of distance functions that allows to build any diagram having such bisectors. We use a linearization technique to study this question.

2.5.1 Abstract Diagrams

Voronoi diagrams have been defined (see Sect. 2.2) as the minimization diagram of a finite set of continuous functions $\{\delta_1, \ldots, \delta_n\}$. It is convenient to interpret each δ_i as the distance function to an abstract object o_i, $i = 1, \ldots, n$. We define the bisector of two objects o_i and o_j of $\mathcal{O} = \{o_1, \ldots, o_n\}$ as

$$b_{ij} = \{x \in \mathbb{R}^d, \delta_i(x) = \delta_j(x)\}.$$

The bisector b_{ij} subdivides $\mathbb{R}^d$ into two open regions: one, b_{ij}^i, consisting of the points of $\mathbb{R}^d$ that are closer to o_i than to o_j, and the other one, b_{ij}^j, consisting of the points of $\mathbb{R}^d$ that are closer to o_j than to o_i. We can then define the Voronoi region of o_i as the intersection of the regions b_{ij}^i for all $j \neq i$. The union of the closures of these Voronoi regions covers $\mathbb{R}^d$. Furthermore, if we assume that the bisectors are $(d-1)$-manifolds, the Voronoi regions then have disjoint interiors and we can define the closed region associated to b_{ij}^i as $\bar{b}_{ij}^i = b_{ij}^i \cup b_{ij}$.

In a way similar to Klein [230], we now define diagrams in terms of bisectors instead of distance functions. Let $B = \{b_{ij}, i \neq j\}$ be a set of closed $(d-1)$-manifolds without boundary. We always assume in the following that $b_{ij} = b_{ji}$ for all $i \neq j$. We assume further that, for all distinct i, j, k, the following incidence condition (I.C.) holds:

$$b_{ij} \cap b_{jk} = b_{jk} \cap b_{ki} \quad (I.C.)$$

This incidence condition is obviously needed for B to be the set of bisectors of some distance functions.

By Jordan's theorem, each element of B subdivides $\mathbb{R}^d$ into at least two connected components and crossing a bisector b_{ij} implies moving into another

connected component of $\mathbb{R}^d \setminus b_{ij}$. Hence, once a connected component of $\mathbb{R}^d \setminus b_{ij}$ is declared to belong to i, the assignments of all the other connected components of $\mathbb{R}^d \setminus b_{ij}$ to i or j are determined.

Given a set of bisectors $B = \{b_{ij}, i \neq j\}$, an *assignment* on B associates to each connected component of $\mathbb{R}^d \setminus b_{ij}$ a label i or j so that two adjacent connected components have different labels.

Once an assignment on B is defined, the elements of B are called *oriented bisectors*.

Given B, let us now consider such an assignment and study whether it may derive from some distance functions. In other words, we want to know whether there exists a set $\Delta = \{\delta_1, \ldots, \delta_n\}$ of distance functions such that

1. the set of bisectors of Δ is B;
2. for all $i \neq j$, a connected component C of $\mathbb{R}^d \setminus b_{ij}$ is labeled by i if and only if
$$\forall x \in C, \delta_i(x) \leq \delta_j(x).$$

We define the *region* of object o_i as $\cap_{j \neq i} \bar{b}^i_{ij}$.

A necessary condition for the considered assignment to derive from some distance functions is that the regions of any subdiagram cover $\mathbb{R}^d$. We call this condition the assignment condition (A.C.):

$$\forall I \subset \{1, \ldots, n\}, \cup_{i \in I} \cap_{j \in I \setminus \{i\}} \bar{b}^i_{ij} = \mathbb{R}^d \quad (A.C.)$$

Given a set of bisectors $B = \{b_{ij}, i \neq j\}$ and an assignment satisfying I.C. and A.C., the *abstract diagram* of $\mathcal{O}$ is the subdivision of $\mathbb{R}^d$ consisting of the regions of the objects of $\mathcal{O}$ and of their faces. The name *abstract Voronoi diagram* was coined by Klein [230], referring to similar objects in the plane.

For any set of distance functions δ_i, we can define the corresponding set of oriented bisectors. Obviously, I.C. and A.C. are satisfied and the abstract diagram defined by this set is exactly the minimization diagram for the distance functions δ_i. Hence any Voronoi diagram allows us to define a corresponding abstract diagram. Let us now prove the converse: any abstract diagram can be constructed as a Voronoi diagram.

Specifically, we prove that I.C. and A.C. are sufficient conditions for an abstract diagram to be the minimization diagram of some distance functions, thus proving the equivalence between abstract diagrams and Voronoi diagrams. We need the following technical lemmas.

Lemma 2. *The assignment condition implies that for any distinct i, j, k, we have*
$$b^j_{ij} \cap b^k_{jk} \cap b^i_{ki} = \emptyset.$$

Proof. A.C. implies that $\mathbb{R}^d = \cup_{1 \leq i \leq n} \cap_{j \neq i} \bar{b}^i_{ij} \subset \bar{b}^i_{ij} \cup \bar{b}^j_{jk} \cup \bar{b}^k_{ki}$. Hence, $\bar{b}^i_{ij} \cup \bar{b}^j_{jk} \cup \bar{b}^k_{ki} = \mathbb{R}^d$. Taking the complementary sets, we obtain $b^j_{ij} \cap b^k_{jk} \cap b^i_{ki} = \emptyset$.

Lemma 3. *For any distinct i, j, k, we have*

$$b_{ij} \cap b^k_{jk} \subset b^k_{ik} \quad \textit{and} \quad b_{ij} \cap \bar{b}^k_{jk} \subset \bar{b}^k_{ik} \tag{2.3}$$

$$b_{ij} \cap b^j_{jk} \subset b^i_{ik} \quad \textit{and} \quad b_{ij} \cap \bar{b}^j_{jk} \subset \bar{b}^i_{ik} \tag{2.4}$$

Proof. Let us first prove that $b_{ij} \cap b^k_{jk} \subset b^k_{ik}$:

Consider $x \in b_{ij} \cap b^k_{jk}$. Assume, for a contradiction, that $x \notin b^k_{ik}$. It follows that $x \in \bar{b}^i_{ik}$, but x cannot lie on b_{ik}, because this would imply that $x \in b_{ik} \cap b_{ij}$, which does not intersect b^k_{jk}. Hence, $x \in b^i_{ik}$ and therefore, $x \in b_{ij} \cap b^k_{jk} \cap b^i_{ik}$. We can then find x' in the neighborhood of x such that $x' \in b^j_{ij} \cap b^k_{jk} \cap b^i_{ki}$, contradicting Lemma 2.

Let us now prove that $b_{ij} \cap \bar{b}^k_{jk} \subset \bar{b}^k_{ik}$. We have proved the inclusion for $b_{ij} \cap b^k_{jk}$. It remains to prove that $b_{ij} \cap b_{jk} \subset \bar{b}^k_{ik}$ which is trivially true, by I.C.

The two other inclusions are proved in a similar way.

We can now prove a lemma stating a transitivity relation:

Lemma 4. *For any distinct i, j, k, we have $b^i_{ij} \cap b^j_{jk} \subset b^i_{ik}$.*

Proof. Let $x \in b^i_{ij} \cap b^j_{jk}$. Assume, for a contradiction, that $x \notin b^i_{ik}$. If $x \in b^k_{ik}$, we have $b^i_{ij} \cap b^j_{jk} \cap b^k_{ik} \neq \emptyset$, contradicting Lemma 2. Therefore, x has to belong to b_{ik}, which implies that $x \in b^i_{ij} \cap b_{ik} \subset b^k_{kj}$ by Lemma 3. This contradicts $x \in b^j_{jk}$. We deduce that $x \in b^i_{ik}$, as needed.

The following lemma states that at most two assignments are likely to derive from some Voronoi diagram.

Lemma 5. *For a given set B satisfying I.C. and assuming that we never have $b_{ij} \subset b_{ik}$ for $j \neq k$, there are at most two ways of labeling the connected components of each $\mathbb{R}^d \setminus b_{ij}$ as b^i_{ij} and b^j_{ij} such that A.C. is verified.*

Proof. First assume that the sides b^1_{12} and b^2_{12} have been assigned. Consider now the labeling of the sides of b_{1i}, for some $i > 2$: let x be a point in the non empty set $b_{2i} \setminus b_{12}$. First assume that $x \in b^1_{12}$. Lemma 3 then implies that $x \in b^1_{1i}$. Conversely, if $x \in b^2_{12}$, $x \in b^i_{1i}$. In both cases, the assignment of the sides of b_{1i} is determined.

All other assignments are determined in a similar way. One can easily see that reversing the sides of b_{12} reverses all the assignments. Thus, we have at most two possible global assignments.

Theorem 12. *Given a set of bisectors $B = \{b_{ij}, 1 \leq i \neq j \leq n\}$ that satisfies the incidence condition (I.C.) and an assignment that satisfies the assignment condition (A.C.), there exists a set of distance functions $\{\delta_i, 1 \leq i \leq n\}$ defining the same bisectors and assignments.*

Proof. Let δ_1 be any real continuous function over $\mathbb{R}^d$. Let $j > 1$ and assume the following induction property: for all $i < j$, the functions δ_i have already been constructed so that

$$\forall i, i' < j, \quad \delta_i(x) \leq \delta_{i'}(x) \Leftrightarrow x \in \bar{b}^i_{ii'}.$$

Let us build δ_j. We consider the arrangement of all bisectors b_{ij}, for $i < j$: for each $I \subset J = \{1, \ldots, j-1\}$, we define $V_I = (\cap_{i \in I} \bar{b}^i_{ij}) \cap (\cap_{k \in J \setminus I} \bar{b}^j_{jk})$. The set V_I is a non necessarily connected region of the arrangement where we need $\delta_j > \delta_i$ if $i \in I$ and $\delta_j < \delta_i$ if $i \in J \setminus I$. This leads us to the following construction.

The interior of V_I is $\operatorname{int} V_I = (\cap_{i \in I} b^i_{ij}) \cap (\cap_{k \in J \setminus I} b^j_{jk})$. Lemma 4 and the induction hypothesis imply that

$$\forall i \in I, \forall k \in J \setminus I, \forall x \in \operatorname{int} V_I, \delta_i(x) < \delta_k(x).$$

In particular, if we define $\nu_I = \min_{k \in J \setminus I} \delta_k$ and $\mu_I = \max_{i \in I} \delta_i$ on V_I, we have $\mu_I < \nu_I$ on $\operatorname{int} V_I$.

Let us now consider some point x on the boundary of V_I. We distinguish two cases. We can first assume that $x \in b_{ij}$ for some $i \in I$. Then, by Lemma 3, for any $i' \in I \setminus \{i\}$, $x \in b_{ij} \cap \bar{b}^{i'}_{i'j} \subset \bar{b}^{i'}_{i'i}$ so that $\delta_{i'}(x) \leq \delta_i(x)$. It follows that $\mu_I(x) = \delta_i(x)$.

Consider now the case when $x \in \partial V_I \cap b_{jk}$ with $k \in J \setminus I$, we have $\nu_I(x) = \delta_k(x)$. Finally, if $x \in \partial V_I \cap b_{ij} \cap b_{jk}$ with $i \in I$ and $k \in J \setminus I$, we have $\mu_I(x) = \delta_i(x)$ and $\nu_I(x) = \delta_k(x)$. By the induction hypothesis, $\delta_i(x) = \delta_k(x)$, which implies that $\mu_I(x) = \nu_I(x)$.

It follows that we can define a continuous function ρ on ∂V_I in the following way:

$$\begin{aligned} \rho_I(x) &= \mu_I(x) \text{ if } \exists i \in I, x \in b_{ij} \\ &= \nu_I(x) \text{ if } \exists k \in J \setminus I, x \in b_{jk} \end{aligned}$$

Furthermore, on $\partial V_I \cap b_{ij} = \partial V_{I \setminus \{i\}} \cap b_{ij}$, if $i \in I$, we have

$$\rho_I(x) = \mu_I(x) = \nu_{I \setminus \{i\}}(x) = \rho_{I \setminus \{i\}}(x). \tag{2.5}$$

The definitions of the ρ_I are therefore consistent, and we can now use these functions to prove that the following definition of δ_j satisfies the induction property.

Finally, we require δ_j to be any continuous function verifying

$$\mu_I < \delta_j < \nu_I$$

on each $\operatorname{int} V_I$. By continuity of δ_j, we deduce from 2.5 that if $x \in \partial V_I \cap b_{jk} = \partial V_{I \setminus \{i\}} \cap b_{ij}$ with $k \in J \setminus I$, we have $\rho_I(x) = \mu_I(x) = \nu_{I \setminus \{i\}}(x) = \rho_{I \setminus \{i\}}(x) = \delta_j(x)$.

It follows that on each V_I, for all $i < j$, $\delta_i(x) < \delta_j(x)$ iff $x \in b^i_{ij}$ and $\delta_i(x) = \delta_j(x)$ iff $x \in b_{ij}$. The induction follows.

One can prove that, in the proof of Lemma 5, the assignment we build satisfies the consequences of A.C. stated in Lemmas 2, 3 and 4. The proof of Theorem 12 does not need A.C. but only the consequences of A.C. stated in those three lemmas. It follows that any of the two possible assignments determined in the proof of Lemma 5 allows the construction of distance functions, as in Theorem 12, which implies that A.C. is indeed verified. We thus obtain a stronger version of Lemma 5.

Lemma 6. *For a given set B satisfying I.C. and assuming that we never have $b_{ij} \subset b_{ik}$ for $j \neq k$, there are* exactly *two ways of labeling the connected components of each $\mathbb{R}^d \setminus b_{ij}$ as b_{ij}^i and b_{ij}^j such that A.C. is verified.*

Theorem 12 proves the equivalence between Voronoi diagrams and abstract diagrams by constructing a suitable set of distance functions. In the case of affine bisectors, the following result of Aurenhammer [35] allows us to choose the distance functions in a smaller class than the class of continuous functions.

Theorem 13. *Any abstract diagram of $\mathbb{R}^d$ with affine bisectors is identical to the power diagram of some set of spheres of $\mathbb{R}^d$.*

Proof. In this proof, we first assume that the affine bisectors are in general position, i.e. four of them cannot have a common subspace of co-dimension 2: the general result easily follows by passing to the limit.

Let $B = \{b_{ij}, 1 \leq i \neq j \leq n\}$ be such a set. We identify $\mathbb{R}^d$ with the hyperplane $x_{d+1} = 0$ of $\mathbb{R}^{d+1}$. Assume that we can find a set of hyperplanes $\{H_i, 1 \leq i \leq n\}$ of $\mathbb{R}^{d+1}$ such that the intersection $H_i \cap H_j$ projects onto b_{ij}. Sect. 2.3 then shows that the power diagram of the set of spheres $\{\sigma_i, 1 \leq i \leq n\}$ obtained by projecting the intersection of paraboloid $\mathcal{Q}$ with each H_i onto $\mathbb{R}^d$ admits B as its set of bisectors[2] (see Fig. 2.5).

Let us now construct such a set of hyperplanes, before considering the question of the assignment condition.

Let H_1 and H_2 be two non-vertical hyperplanes of $\mathbb{R}^d$ such that $H_1 \cap H_2$ projects vertically onto b_{12}. We now define the H_i for $i > 2$: let Δ_i^1 be the maximal subspace of H_1 that projects onto b_{1i} and let Δ_i^2 be the maximal subspace of H_2 that projects onto b_{2i}. Both Δ_i^1 and Δ_i^2 have dimension $d-1$. I.C. implies that $b_{12} \cap b_{2i} \cap b_{i1}$ has co-dimension 2 in $\mathbb{R}^d$. Thus $\Delta_i^1 \cap \Delta_i^2$, its preimage on H_1 (or H_2) by the vertical projection, has the same dimension $d-2$. This proves that Δ_i^1 and Δ_i^2 span a hyperplane H_i of $\mathbb{R}^{d+1}$. The fact that $H_i \neq H_1$ and $H_i \neq H_2$ easily follows from the general position assumption.

We still have to prove that $H_i \cap H_j$ projects onto b_{ij} for $i \neq j > 2$. I.C. ensures that the projection of $H_i \cap H_j$ contains the projection of $H_i \cap H_j \cap H_1$ and the projection of $H_i \cap H_j \cap H_2$, which are known to be $b_{ij} \cap b_{1i}$ and $b_{ij} \cap b_{2i}$, by construction. The general position assumption implies that there is only

[2] We may translate the hyperplanes vertically in order to have a non-empty intersection, or we may consider imaginary spheres with negative squared radii.

one hyperplane of $\mathbb{R}^d$, namely b_{ij}, containing both $b_{ij} \cap b_{1i}$ and $b_{ij} \cap b_{2i}$. This is the projection of $H_i \cap H_j$.

As we have seen, building this set of hyperplanes of $\mathbb{R}^{d+1}$ amounts to building a family of spheres whose power diagram admits B as its set of bisectors. At the beginning of the construction, while choosing H_1 and H_2, we may obtain any of the two possible labellings of the sides of b_{12}. Since there is no other degree of freedom, this choice determines all the assignments. Lemma 5 shows that there are at most two possible assignments satisfying A.C., which proves we can build a set of spheres satisfying any of the possible assignments. The result follows.

Exercise 22. Consider the diagram obtained from the Euclidean Voronoi diagram of n points by taking the other assignment. Characterize a region in this diagram in terms of distances to the points and make a link with Exercise 3.

2.5.2 Inverse Problem

We now assume that each bisector is defined as the zero-set of some real-valued function over $\mathbb{R}^d$, called a *bisector-function* in the following. Let us denote by B the set of bisector-functions. By convention, for any bisector-function β_{ij}, we assume that

$$b^i_{ij} = \{x \in \mathbb{R}^d : \beta_{ij}(x) < 0\} \text{ and } b^j_{ij} = \{x \in \mathbb{R}^d : \beta_{ij}(x) > 0\}.$$

We now define an algebraic equivalent of the incidence relation in terms of pencil of functions: we say that B satisfies the linear combination condition (L.C.C.) if, for any distinct i, j, k, β_{ki} belongs to the pencil defined by β_{ij} and β_{jk}, i.e.

$$\exists(\lambda, \mu) \in \mathbb{R}^2 \quad \beta_{ki} = \lambda\beta_{ij} + \mu\beta_{jk} \quad (L.C.C.)$$

Note that L.C.C. implies I.C. and that in the case of affine bisectors L.C.C. is equivalent to I.C. Furthermore, it should be noted that, in the case of Voronoi diagrams, the bisector-functions defined as $\beta_{ij} = \delta_i - \delta_j$ obviously satisfy L.C.C.

We now prove that we can view diagrams satisfying L.C.C. as diagrams that can be linearized.

Definition 1. *A diagram $\mathcal{D}$ of n objects in some space E is said to be a pullback of a diagram $\mathcal{D}'$ of m objects in space F by a function $\phi : E \to F$ if $m = n$ and if, for any distinct i, j, we have*

$$b^i_{ij} = \phi^{-1}(c^i_{ij})$$

where b^i_{ij} denotes the set of points closer to i than to j in $\mathcal{D}$ and c^i_{ij} denotes the set of points closer to i than to j in $\mathcal{D}'$.

Theorem 14. *Let $B = \{\beta_{ij}\}$ be a set of real-valued bisector-functions over $\mathbb{R}^d$ satisfying L.C.C. and A.C. Let V be any vector space of real functions over $\mathbb{R}^d$ that contains B and constant functions.*

If N is the dimension of V, the diagram defined by B is the pullback by some continuous function of an affine diagram in dimension $N-1$.

More explicitly, there exist a set $C = \{\psi_{ij} \cdot X + c_{ij}\}$ of oriented affine hyperplanes of $\mathbb{R}^{N-1}$ satisfying I.C. and A.C. and a continuous function ϕ: $\mathbb{R}^d \to \mathbb{R}^{N-1}$ such that for all $i \neq j$,

$$\bar{b}^i_{ij} = \{x \in \mathbb{R}^d, \beta_{ij}(x) \leq 0\} = \phi^{-1}\{y \in \mathbb{R}^{N-1}, \psi_{ij}(x) \leq c_{ij}\}.$$

Proof. Let $(\gamma_0, \ldots, \gamma_{N-1})$ be a basis of V such that γ_0 is the constant function equal to 1.

Consider the evaluation application,

$$\phi : x \in \mathbb{R}^d \mapsto (\gamma_1(x), \ldots, \gamma_{N-1}(x)) \in \mathbb{R}^{N-1}.$$

If point x belongs to some b^i_{ij}, we have $\beta_{ij}(x) < 0$. Furthermore, there exists real coefficients $\lambda^0_{ij}, \ldots, \lambda^{N-1}_{ij}$ such that $\beta_{ij} = \sum_{k=0}^{N-1} \lambda^k_{ij}\gamma_k$. The image $\phi(x)$ of x thus belongs to the affine half-space B^i_{ij} of $\mathbb{R}^{N-1}$ of equation

$$\sum_{k=1}^{N-1} \lambda^k_{ij} X_k < -\lambda^0_{ij}.$$

In this way, we can define all the affine half-spaces B^i_{ij} of $\mathbb{R}^{N-1}$ for $i \neq j$: B_{ij} is an oriented affine hyperplane with normal vector $(\lambda^1_{ij}, \ldots, \lambda^{N-1}_{ij})$ and constant term λ^0_{ij}. Plainly, L.C.C. on the β_{ij} translates into I.C. on the B_{ij}, and we have

$$\bar{b}^i_{ij} = \{x \in \mathbb{R}^d, \beta_{ij}(x) \leq 0\} = \phi^{-1}\{y \in \mathbb{R}^{N-1}, B_{ij}(x) \leq -\lambda^n_{ij}\} \tag{2.6}$$

Finally, let us prove that A.C. is also satisfied. Lemma 6 states that the B_{ij} have exactly two inverse assignments satisfying A.C. Furthermore, Equation 2.6 implies that any of these two assignments defines an assignment for the b_{ij} that also satisfies A.C. It follows that if the current assignment did not satisfy A.C., there would be more than two assignments for the b_{ij} that satisfy A.C. This proves that A.C. is also satisfied by the B_{ij} and concludes the proof.

We can now use Theorem 13 and specialize Theorem 14 to the specific case of diagrams whose bisectors are hyperspheres or hyperquadrics, or, more generally, to the case of diagrams whose class of bisectors spans a finite dimensional vector space.

Theorem 15. *Any abstract diagram of $\mathbb{R}^d$ with spherical bisectors such that the corresponding degree 2 polynomials satisfy L.C.C. is a Möbius diagram.*

Proof. Since the spherical bisectors satisfy L.C.C., we can apply Theorem 14 and Theorem 13. Function ϕ of Theorem 14 is simply the lifting mapping $x \mapsto (x, x^2)$, and we know from Theorem 13 that our diagram can be obtained as a power diagram pulled-back by ϕ. That is to say $\delta_i(x) = \Sigma_i(\phi(x))$, where Σ_i is a hypersphere in $\mathbb{R}^{d+1}$.

Another way to state this transformation is to consider the diagram with spherical bisectors in $\mathbb{R}^d$ as the projection by ϕ^{-1} of the restriction of the power diagram of the hyperspheres Σ_i to the paraboloid $\phi(\mathbb{R}^d) \subset \mathbb{R}^{d+1}$ of equation $x_{d+1} = x^2$.

Assume that the center of Σ_j is $(u_1^j, \ldots, u_{d+1}^j)$, and that the squared radius of Σ_j is w^j. We denote by Σ_j the power to Σ_j. Distance δ_j can be expressed in terms of these parameters:

$$\delta_j(x) = \Sigma_j(\phi(x)) = \sum_{1 \le i \le d} (x_i - u_i^j)^2 + ((\sum_{1 \le i \le d} x_i^2) - u_{d+1}^j)^2 - w^j.$$

Subtracting from each δ_j the same term $(\sum_{1 \le i \le D} x_i^2)^2$ leads to a new set of distance functions that define the same minimization diagram as the δ_j. In this way, we obtain new distance functions which are exactly the ones defining Möbius diagrams.

This proves that any diagram whose bisectors are hyperspheres can be constructed as a Möbius diagram.

The proof of the following theorem is similar to the previous one:

Theorem 16. *Any abstract diagram of $\mathbb{R}^d$ with quadratic bisectors such that the corresponding degree 2 polynomials satisfy L.C.C. is an anisotropic Voronoi diagram.*

Exercise 23. Explain why, in Theorem 15, it is important to specify which bisector-functions satisfy L.C.C. instead of mentioning only the bisectors (Hint: Theorem 12 implies that there always exist some bisector-functions with the same zero-sets that satisfy L.C.C.)

2.6 Incremental Voronoi Algorithms

Incremental constructions consist in adding the objects one by one in the Voronoi diagram, updating the diagram at each insertion. Incremental algorithms are well known and highly popular for constructing Euclidean Voronoi diagrams of points and power diagrams of spheres in any dimension. Because the whole diagram can have to be modified at each insertion, incremental algorithms have a poor worst-case complexity. However most of the insertions

result only in local modifications and the worst-case complexity does not reflect the actual complexity of the algorithm in most practical situations. To provide more realistic results, incremental constructions are analyzed in the randomized framework. This framework makes no assumption on the input object set but analyzes the expected complexity of the algorithm assuming that the objects are inserted in random order, each ordering sequence being equally likely. The following theorem, whose proof can be found in many textbooks (see e.g. [67]) recalls that state-of-the-art incremental constructions of Voronoi diagrams of points and power diagrams have an optimal randomized complexity.

Theorem 17. *The Euclidean Voronoi diagram of n points in $\mathbb{R}^d$ and the power diagram of n spheres in $\mathbb{R}^d$ can be constructed by an incremental algorithm in randomized time $O\left(n \log n + n^{\lceil \frac{d+1}{2} \rceil}\right)$.*

Owing to the linearization techniques of Sect. 2.5, this theorem yields complexity bounds for the construction of linearizable diagrams such as Möbius, anisotropic or Apollonius diagrams. Incremental constructions also apply to the construction of Voronoi diagrams for which no linearization scheme exists. This is for instance the case for the 2-dimensional Euclidean Voronoi diagrams of line segments. The efficiency of the incremental approach merely relies on the fact that the cells of the diagram are simply connected and that the 1-skeleton of the diagram, (i. e. the union of its edges and vertices) is a connected set. Unfortunately, these two conditions are seldom met except for planar Euclidean diagrams. Let us take Apollonius diagrams as an illustration. Each cell of an Apollonius diagram is star shaped with respect to the center of the associated sphere and is thus simply connected. In the planar case, Apollonius bisectors are unbounded hyperbolic arcs and the 1-skeleton can easily be made connected by adding a curve at infinity. The added curve can be seen as the bisector separating any input object from an added fictitious object. In 3-dimensional space, the skeleton of Apollonius diagrams is not connected: indeed, we know from Sect. 2.4.3 that the faces of a single cell are in 1-1 correspondence with the faces of a 2-dimensional Möbius diagram and therefore may include isolated loops.

As a consequence, the rest of this section focuses on planar Euclidean diagrams. After some definitions, the section recalls the incremental construction of Voronoi diagrams, outlines the topological conditions under which this approach is efficient and gives some examples. The efficiency of incremental algorithms also greatly relies on the availability of some point location data structure to answer nearest neighbor queries. A general data structure, the Voronoi hierarchy, is described at the end of the section. The last subsection lists the main predicates involved in the incremental construction of Voronoi diagrams.

2.6.1 Planar Euclidean diagrams

To be able to handle planar objects that possibly intersect, the distance functions that we consider in this section are signed Euclidean distance functions, i.e. the distance $\delta_i(x)$ from a point x to an object o_i is:

$$\delta_i(x) = \begin{cases} \min_{y \in \bar{o}_i} \|y - x\|, & \text{if } x \notin o \\ -\min_{y \in \bar{o}_i^c} \|y - x\|, & \text{if } x \in o \end{cases}$$

where $\bar{o}_i$ is the closure of o_i and $\bar{o}_i^c$ the closure of the complement of o_i. Note that the distance used to define Apollonius diagrams matches this definition. Then, given a finite set $\mathcal{O}$ of planar objects and $o_i \in \mathcal{O}$, we define the Voronoi region of o_i as the locus of points closer to o_i than to any other object in $\mathcal{O}$

$$V(o_i) = \{x \in \mathbb{R}^2 : \delta_i(x) \leq \delta_j(x),\ \forall o_j \in \mathcal{O}\}.$$

Voronoi edges are defined as the locus of points equidistant to two objects $\mathcal{O}$ and closer to these two objects than to any other object in $\mathcal{O}$, and Voronoi vertices are the locus of points equidistant to three or more objects and closer to these objects than to any other object in $\mathcal{O}$. The Voronoi diagram $\mathrm{Vor}(\mathcal{O})$ is the planar subdivision induced by the Voronoi regions, edges and vertices.

The incremental construction described below relies on the three following topological properties of the diagram that are assumed to be met for any set of input objects:

1. The diagram is assumed to be a *nice diagram*, i. e. a diagram in which edges and vertices are respectively 0 and 1-dimensional sets.
2. The cells are assumed to be simply connected.
3. The 1-*skeleton* of the diagram is connected.

Owing to Euler formula, Properties 1 and 2 imply that the Voronoi diagram of n objects is a planar map of complexity $O(n)$. Property 3 is generally not granted for any input set. Think for example of a set of points on a line. However, in the planar case, this condition can be easily enforced as soon as Properties 1 and 2 are met. Indeed, if the cells are simply connected, there is no bounded bisector and the 1-*skeleton* can be connected by adding a curve at infinity. The added curve can be seen as the bisector separating any input object from an added fictitious object. The resulting diagram is called the *compactified* version of the diagram.

2.6.2 Incremental Construction

We assume that the Voronoi diagram of any input set we consider is a nice diagram with simply connected cells and a connected 1-skeleton. Each step of the incremental construction takes as input the Voronoi diagram $\mathrm{Vor}(\mathcal{O}_{i-1})$ of a current set of objects $\mathcal{O}_{i-1}$ and an object $o_i \notin \mathcal{O}_{i-1}$, and aims to construct

the Voronoi diagram $\mathrm{Vor}(\mathcal{O}_i)$ of the set $\mathcal{O}_i = \mathcal{O}_{i-1} \cup \{o_i\}$. In the following, we note $V(o, \mathcal{O}_{i-1})$ the region of an object o in the diagram $\mathrm{Vor}(\mathcal{O}_{i-1})$ and $V(o, \mathcal{O}_i)$ the region of o in $\mathrm{Vor}(\mathcal{O}_i)$.

We note $\mathrm{Skel}(\mathcal{O}_{i-1})$ the 1-skeleton of $\mathrm{Vor}(\mathcal{O}_{i-1})$. Let x be a point of $\mathrm{Skel}(\mathcal{O}_{i-1})$. We note $\mathcal{N}(x, \mathcal{O}_{i-1})$ the nearest neighbors of x in $\mathcal{O}_{i-1}$, i.e. the subset of objects of $\mathcal{O}_{i-1}$ that are closest to x. The point x is said to be in conflict with o_i if x is closer to o_i than to $\mathcal{N}(x, \mathcal{O}_{i-1})$. Hence, the part of the skeleton that conflicts with o_i, called the *conflict skeleton* for short, is exactly the intersection of the skeleton $\mathrm{Skel}(\mathcal{O}_{i-1})$ with the region of o_i in $\mathrm{Vor}(\mathcal{O}_i)$. See Fig. 2.11.

The conflict skeleton is a subgraph of $\mathrm{Skel}(\mathcal{O}_{i-1})$ and the endpoints of this subgraph are the vertices of the region $V(o_i, \mathcal{O}_i)$. If $V(o_i, \mathcal{O}_i)$ is not empty, the conflict skeleton is not empty either. Indeed, an empty conflict skeleton would imply that $V(o_i, \mathcal{O}_i)$ is included in a single region $V(o, \mathcal{O}_{i-1})$ of the diagram $\mathrm{Vor}(\mathcal{O}_{i-1})$ and the region $V(o, \mathcal{O}_i)$ would not be simply connected. Furthermore, the following lemma, due to Klein et al. [231], proves that the conflict skeleton is a connected subgraph of $\mathrm{Skel}(\mathcal{O}_{i-1})$.

Lemma 7. *If, for any input set, the Voronoi diagram is a nice diagram with simply connected regions and a connected* 1*-skeleton, the conflict skeleton of an additional object is connected.*

Proof. We use the above notation and assume, for a contradiction, that the conflict skeleton of o_i, which is $\mathrm{Skel}(\mathcal{O}_{i-1}) \cap V(o_i, \mathcal{O}_i)$, consists of several disjoint connected components $Sk_1, Sk_2, \ldots, Sk_\ell$. Each connected component Sk_j has to intersect the boundary of the new region $V(o_i, \mathcal{O}_i)$, otherwise $\mathrm{Skel}(\mathcal{O}_{i-1})$ would not be connected, a contradiction. Then, if $\ell \geq 2$, there exists a path $\mathcal{C}$ in $V(o_i, \mathcal{O}_i)$ connecting two points x and y on the boundary of $V(o_i, \mathcal{O}_i)$ and separating Sk_1 and Sk_2, see Fig. 2.10. The path $\mathcal{C}$ does not intersect $\mathrm{Skel}(\mathcal{O}_{i-1})$ and is therefore included in the region $V(o, \mathcal{O}_{i-1})$ of some object o of $\mathcal{O}_{i-1}$. Since, points arbitrarily close to x and y but outside $V(o_i, \mathcal{O}_i)$ belong to $V(o, \mathcal{O}_{i-1})$, x and y can be joined by a simple path $\mathcal{D}$ in $V(o, \mathcal{O}_i) \subset V(o, \mathcal{O}_{i-1})$. The simple closed curve $\mathcal{C} \cup \mathcal{D}$ is contained in $V(o, \mathcal{O}_{i-1})$ and encloses Sk_1 or Sk_2, which contradicts the fact that $\mathrm{Skel}(\mathcal{O}_{i-1})$ is connected.

Once the conflict skeleton is known, the Voronoi diagram $\mathrm{Vor}(\mathcal{O}_{i-1})$ can be updated, leading to $\mathrm{Vor}(\mathcal{O}_i)$. This is done by Procedure 3.

Procedure 4 describes a step of the incremental construction.

In the sequel, the incremental construction is analyzed in the randomized setting. It is assumed that each object has constant complexity, which implies that each operation involving a constant number of objects is performed in constant time. Because the conflict skeleton is connected, Substep 2 can be performed by traversing the graph $\mathrm{Skel}(\mathcal{O}_{i-1})$ in time proportional to the number of edges involved in the conflict skeleton. These edges will be either deleted or shortened in the new diagram. Substep 3 takes time proportional

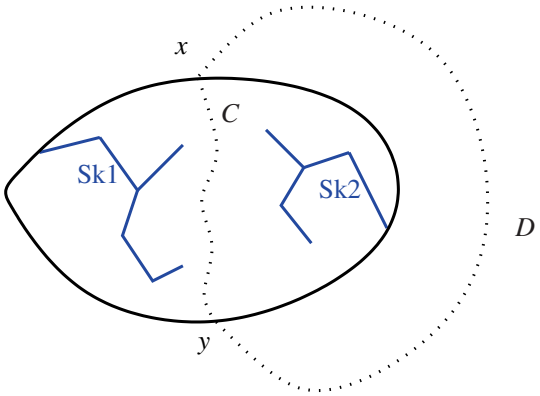

Fig. 2.10. For the proof of Lemma 7

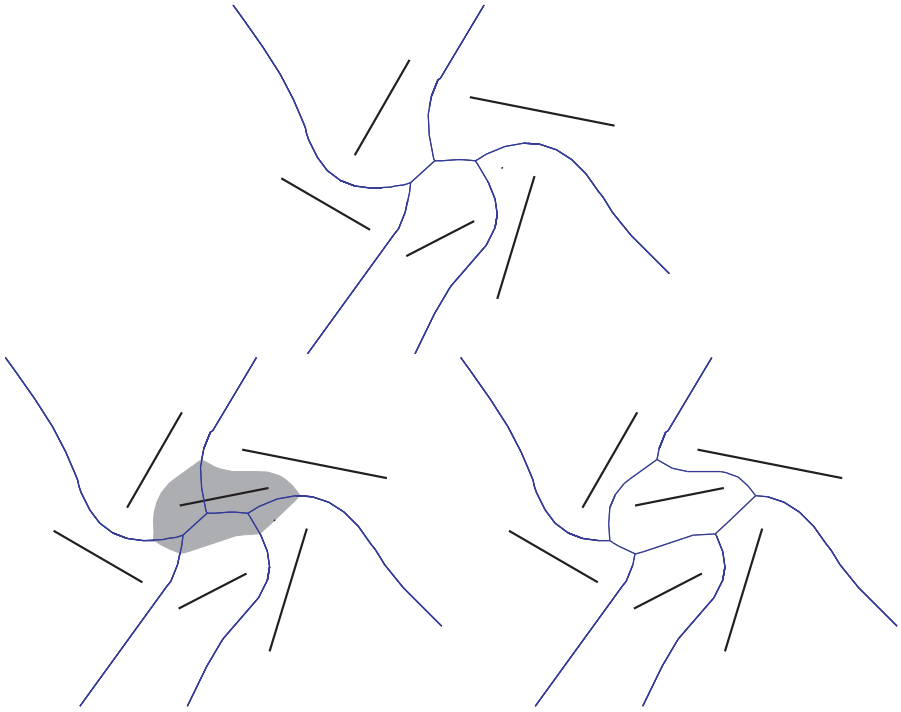

Fig. 2.11. Incremental construction of the Voronoi diagram of disjoint line segments

Procedure 3 Updating the Voronoi diagram

INPUT: $\mathrm{Vor}(\mathcal{O}_{i-1}), \mathrm{Skel}(\mathcal{O}_{i-1})$

1. Create a new vertex at each endpoint of the conflict skeleton;
2. Remove vertices, edges and portions of edges that belong to $\mathrm{Skel}(\mathcal{O}_{i-1}) \cap V(o_i, \mathcal{O}_i)$;
3. Connect the new vertices s as to form the boundary of the new region.

OUTPUT: $\mathrm{Vor}(\mathcal{O}_i)$

Procedure 4 A step of the incremental algorithm

INPUT: $\mathrm{Vor}(\mathcal{O}_{i-1})$ and a new object o_i

1. Find a first point x of $\mathrm{Skel}(\mathcal{O}_{i-1})$ in conflict with o_i;
2. Compute the whole conflict skeleton;
3. Update $\mathrm{Vor}(\mathcal{O}_{i-1})$ into $\mathrm{Vor}(\mathcal{O}_i)$ using Procedure 3;
4. Update the location data structure.

OUTPUT: $\mathrm{Vor}(\mathcal{O}_i)$

to the number of edges involved in the conflict skeleton plus the number of edges of $V(o_i, \mathcal{O}_i)$. The latter are the new edges. Hence Substeps 2 and 3 take time proportional to the number of changes in the 1-skeleton. Because each edge in the skeleton is defined by four objects and because the complexity of the Voronoi diagram of n objects is $O(n)$, a standard probabilistic analysis (see e.g. [67]) shows that the expected number of changes at each step of the incremental algorithm is $O(1)$. The overall randomized complexity of the algorithm is $O(n)$.

The costs of Substeps 1 and 4 depend of course on the type of the input objects and of the location data structure. In Sect. 2.6.3, we described a location data structure, called the Voronoi hierarchy, that can be used in the case of disjoint convex objects. The Voronoi hierarchy allows to detect a first conflict in randomized time $O(\log^2 n)$. At each step, the data structure is updated in time $O(m \log^2 n)$ where m is the number of changes in the diagram. Because the expected number of changes at each step is $O(1)$, the expected cost for updating the hierarchy is $O(\log^2 n)$. This yields the following theorem.

Theorem 18. *The incremental construction of the planar Euclidean Voronoi diagram of n disjoint convex objects with constant complexity takes $O(n \log^2 n)$ expected time.*

Note that the incremental construction described here is *on-line*, meaning that the algorithm does not need to know the whole set of objects right from the beginning. If the whole set of objects is known in advance, localization can be made easy and the maintenance of a location data structure is no longer required [20]. The idea consists in picking one witness point inside each object and in building first the Voronoi diagram of witness points. In a second

phase, each witness point is replaced in turn by the corresponding object. When replacing a witness point by the corresponding object, any point on the boundary of the cell of the witness point belongs to the conflict skeleton. The algorithm is no longer on-line but its randomized complexity is reduced to $O(n \log n)$.

Voronoi Diagrams of Line Segments

The above incremental construction applies to the Voronoi diagram of disjoint line segments. Indeed, in the case of disjoint line segments, the bisector curves are unbounded simple curves, each composed of at most seven line segments and parabolic arcs (see e.g. [67, Chap. 19]). Hence, Voronoi vertices and edges are respectively 0 and 1-dimensional sets. Furthermore, each region in the diagram is weakly star shaped with respect to its generating segment, meaning that the segment joining any point in the region to its closest point on the associated segment is included in the region. It follows that Voronoi regions are simply connected.

If the segments are allowed to share endpoints, the Voronoi diagram exhibits 2-dimensional Voronoi edges, hence violating the definition of nice Voronoi diagrams. A way to circumvent this problem consists in considering that each segment is composed of three distinct objects: the two endpoints and the open segment. If the two endpoints of a segment are inserted in the diagram prior to the open segment, the incremental construction encounters no 2-dimensional bisecting region and the algorithm presented above can be used.

Voronoi Diagrams of Curved Segments

Voronoi diagrams of disjoint curved segments have been studied by Alt, Cheong and Vigneron [20]. Alt, Cheong and Vigneron introduce the notion of harmless curved segments defined as follows. A curved segment is said to be *convex* when the region bounded by the curved segment and the line segment joining its endpoints is convex. A *spiral* arc is a convex curved segment with monotonously increasing curvature. A *harmless* curved segment is either a line segment or a circular arc or a spiral arc. It can be shown that, if the input curved segments are split into harmless sub-segments and if each open curved sub-segment and its two endpoints are considered as three distinct sites, the Voronoi diagram is a nice diagram with simply connected regions. The incremental construction paradigm described above therefore applies.

Voronoi Diagrams of Convex Objects

The case of disjoint smooth convex objects is quite similar to the case of disjoint segments. The bisecting curves between two such objects is a 1-dimensional curve. Furthermore, each Voronoi region is weakly star shaped

with respect to the medial axis of its object, hence simply connected. Therefore, the Voronoi diagram can be built using the incremental algorithm. Note that the Voronoi diagram of disjoint smooth convex objects could also be obtained by applying the incremental algorithm to the curved segments forming the boundaries of the objects. However, this approach requires the subdivision of the boundary of each object into harmless parts and yields a Voronoi diagram which is a refinement of the diagram of the input objects.

If we still assume the objects to be smooth and convex but allow them to intersect, things become more difficult. Karavelas and Yvinec [217] have shown that the Voronoi regions remain simply connected if and only if the objects of $\mathcal{O}$ are pseudo-disks, meaning that the boundaries of any two objects of $\mathcal{O}$ intersect in at most two points. The above incremental algorithm can be adapted to work in this case. However, since the distance is a signed distance, some sites may have an empty region, which makes the algorithm slightly more complicated. Note that this may only happen when some of the objects are included in the union of others. The algorithm has to check that the new object has a non-empty region and must handle the case where the insertion of a new object causes the region of some other object to vanish. Karavelas and Yvinec [217] showed that there is no use to maintain a location data structure in this case because each insertion takes linear time anyway.

The algorithm can be generalized to the case of convex objects with piecewise smooth pseudo-circular boundaries. As in the case of segments, the main problem comes from the fact that sharp corners on the boundaries of objects yield 2-dimensional bisectors. This problem can be handled as in the case of line segments and planar curves by considering each corner as an object on its own.

2.6.3 The Voronoi Hierarchy

The first step when inserting a new object o_i consists in finding one point of the current skeleton $\mathrm{Skel}(\mathcal{O}_{i-1})$ in conflict with o_i. If the objects do not intersect, this is done by searching the object o of $\mathcal{O}_{i-1}$ nearest to a point x of o_i. Indeed, if the objects do not intersect, x belongs to the region of o in $\mathrm{Vor}(\mathcal{O}_{i-1})$ and to the region of o_i in $\mathrm{Vor}(\mathcal{O}_i)$. Therefore o_i has to be a neighbor of o in $\mathrm{Vor}(\mathcal{O}_i)$ and some point on the boundary of $V(o, \mathcal{O}_{i-1})$ is in conflict with o_i. If the objects intersect, things are slightly more complicated but nearest object queries can still be used to find out whether the new object is hidden and, if not, to find a first conflict (see [217]).

Let us describe now how to find the object o of a set $\mathcal{O}$ closest to a query point x. A simple strategy is to perform a walk in the Voronoi diagram $\mathrm{Vor}(\mathcal{O})$. The walk starts at any region of the diagram. When the walk visits the region $V(o)$ of an object o it considers in turn each of the neighboring regions. If one of the neighbors of o, say o', is closer to x than o, the walk steps to the region $V(o')$. If none of the neighbors of o in $\mathrm{Vor}(\mathcal{O})$ is closer to x than o, then o is the object closest to x and the walk ends. Because the distance

between x and the objects of the visited regions is decreasing, the walk cannot loop and is bound to end. However, the walk may visit all the regions before ending. The Voronoi hierarchy [217] is a randomized data structure that makes this strategy more efficient. The Voronoi hierarchy can be considered as a 2-dimensional version of the skip lists introduced by Pugh [291] and generalizes the Delaunay hierarchy described in [115].

For a set of objects $\mathcal{O}$, the Voronoi hierarchy $\mathcal{H}_{\mathcal{V}}(\mathcal{O})$ is a sequence of Voronoi diagrams $\mathrm{Vor}(\Theta_\ell)$, $\ell = 0, \ldots, L$, built for subsets of $\mathcal{O}$ forming a hierarchy, i.e, $\mathcal{O} = \Theta_0 \supseteq \Theta_1 \supseteq \cdots \supseteq \Theta_L$.

The hierarchy $\mathcal{H}_{\mathcal{V}}(\mathcal{O})$ is built together with the Voronoi diagram $\mathrm{Vor}(\mathcal{O})$ according to the following rules:

1. Every object of $\mathcal{O}$ is inserted in $\mathrm{Vor}(\Theta_0) = \mathrm{Vor}(\mathcal{O})$;
2. An object o that has been inserted in $\mathrm{Vor}(\Theta_\ell)$, is inserted in $\mathrm{Vor}(\Theta_{\ell+1})$ with probability β.

To answer nearest object queries, the Voronoi hierarchy works as follows. Let us call θ_ℓ the object of Θ_ℓ closest to the query point x. First, a simple walk is performed in the top-most diagram to find θ_L. Then, at each level $\ell = L-1, \ldots, 0$, a simple walk is performed in $\mathrm{Vor}(\Theta_\ell)$ from $\theta_{\ell+1}$ to θ_ℓ (see Fig. 2.12).

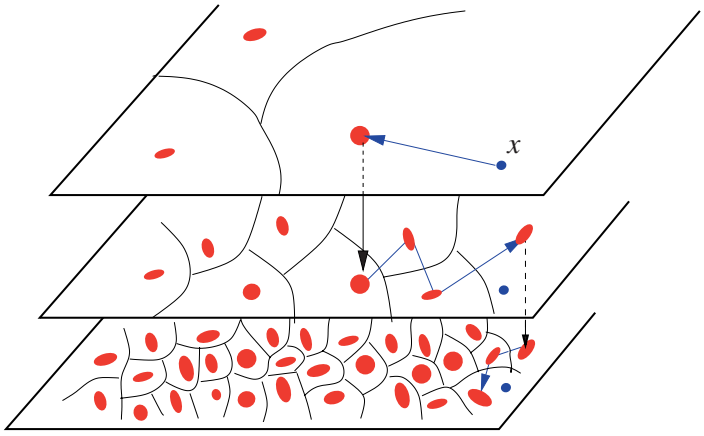

Fig. 2.12. Locating x using the Voronoi hierarchy

It is easy to show that the expected size of $\mathcal{H}_{\mathcal{V}}(\mathcal{O})$ is $O(\frac{n}{1-\beta})$, and that the expected number of levels in $\mathcal{H}_{\mathcal{V}}(\mathcal{O})$ is $O(\log_{1/\beta} n)$. Moreover, the following lemma proves that the expected number of steps performed by the walk at each level is constant.

Lemma 8. *Let x be a point in the plane. Let θ_ℓ (resp. $\theta_{\ell+1}$) be the object closest to x in Θ_ℓ (resp. $\Theta_{\ell+1}$). Then the expected number of Voronoi regions visited during the walk in* $\mathrm{Vor}(\Theta_\ell)$ *from $\theta_{\ell+1}$ to θ_ℓ is $O(1/\beta)$.*

Proof. The objects whose regions are visited at level ℓ are closer to x than $\theta_{\ell+1}$. Consequently, if, among the objects of Θ_ℓ, $\theta_{\ell+1}$ is the k-th closest object to x, the walk in $\mathrm{Vor}(\Theta_\ell)$ performs at most k steps.

Let us show that $\theta_{\ell+1}$ is the k-th closest object to x in Θ_ℓ with probability $\beta(1-\beta)^{k-1}$. Such a case occurs if and only if the two following conditions are satisfied:

1. Object $\theta_{\ell+1}$ has been inserted in $\Theta_{\ell+1}$
2. None of the $k-1$ objects of Θ_ℓ that are closer to x than $\theta_{\ell+1}$ has been inserted in $\Theta_{\ell+1}$.

The first condition occurs with probability β and the second with probability $(1-\beta)^{k-1}$.

Let n_l be the number of objects in Θ_ℓ. The expected number N_ℓ of objects that are visited at level ℓ is bounded as follows:

$$N_\ell \leq \sum_{k=1}^{n_\ell} k(1-\beta)^{k-1}\beta < \beta \sum_{k=1}^{\infty} k(1-\beta)^{k-1} = \frac{1}{\beta},$$

We still have to bound the time spent in each of the visited regions. Let o be the site of a visited region in $\mathrm{Vor}(\Theta_\ell)$. It is not efficient to consider in turn each neighbor o' of o in $\mathrm{Vor}(\Theta_\ell)$ and compare the distances from x to o and o'. Indeed, since the complexity of each region in the Voronoi diagram $\mathrm{Vor}(\Theta_\ell)$ may be $\Omega(n_\ell)$, this would imply that the time spent at each level ℓ of the hierarchy is $O(n)$, yielding a total of $O(n)$ time per insertion. To avoid this cost, a balanced binary tree is attached to each Voronoi region in the Voronoi hierarchy. The tree attached to the region $V_\ell(o)$ of o in $\mathrm{Vor}(\Theta_\ell)$ includes, for each Voronoi vertex v of $V_\ell(o_i)$, the ray $[v_o, v)$ issued from the projection of v onto the boundary ∂o of o that passes through v. The rays are sorted by directions. Similarly, we associate to the query point x the ray $[x_o, x)$ where x_o is the projection of x onto ∂o. When $V_\ell(o)$ is visited, ray $[x_o, x)$ is located in the associated tree and we get the two rays $[v_o, v)$ and $[w_o, w)$ immediately before and after $[x_o, x)$. Let o' be the neighbor of o whose cell is incident to v and w. We compare the distances from x to o and o'. If o' is closer to x than o, the walk steps to o'. Otherwise, we know that o is the object of Θ_ℓ closest to x and the walk halts (see [217] for details). Hence, visiting the Voronoi region of o_i in $\mathrm{Vor}(\Theta_\ell)$ reduces to querying the tree and comparing the distances from x to o and o' which takes $O(\log n_\ell)$ time.

Lemma 9. *Using a hierarchy of Voronoi diagrams, nearest neighbor queries can be answered in expected time $O(\log^2 n)$.*

It has been shown [217] that the expected cost of updating the Voronoi hierarchy when inserting an object is $O(\log^2 n)$.

Exercise 24. Show that the planar Euclidean Voronoi diagram of n points can be computed on line in $O(n \log n)$ time
(Hint: in the case of points, using a Delaunay hierarchy instead of the Voronoi hierarchy, nearest neighbor queries can be answered in $O(\log n)$. Upon insertion the structure is updated in randomized time $O(\log n)$. See [115])

Exercise 25. Show that the planar Euclidean Voronoi diagram of n line segments can be computed on line in $O(n \log n)$ time. See e.g.[67] for a solution.

Exercise 26. Show that the planar euclidean diagram of n disjoint convex objects can be computed using predicates that involve only four objects.

Exercise 27. Provide a detailed description of the predicates of the incremental Voronoi diagram construction and a way to implement them efficiently for various types of simple objects (e.g., line segments, circles). See [219, 218].

Exercise 28 (2D Abstract Voronoi diagrams). Klein et al. [231] have defined abstract Voronoi diagrams in dimension 2 using bisecting curves. Each bisecting curve b_{ij} is assumed to be an infinite curve separating the plane in two regions affected respectively to o_i and o_j and the Voronoi regions are defined as in Sect. 2.5.1. Klein et al. assume that the affectation fulfills the assignment condition but they do not assume the incidence condition. Instead they assume that each pair of bisecting curves intersect in only a finite number of connected components and that the interior of Voronoi regions are path-connected. Show that, under Klein et al. assumptions, the transitivity relation of Lemma 4 is satisfied and that Voronoi regions are simply connected.

Exercise 29. Let $\mathcal{O}$ be a set of planar convex objects that may intersect and may not form a pseudo-circle set. Show that the Voronoi diagram $\mathrm{Vor}(\mathcal{O})$ may exhibit disconnected Voronoi regions. Propose an extension of the incremental algorithm to build the restriction $\mathrm{Vor}(\mathcal{O}) \cap U^c$ of the Voronoi diagram $\mathrm{Vor}(\mathcal{O})$ to the region U^c which is the complement of the union of the objects. The solution can be found in [217].

Exercise 30. Describes the geometric predicates required to implement the incremental algorithm. Provide algebraic expressions for the case of circles and line segments [219, 146].

2.7 Medial Axis

In this section, we introduce the concept of Medial Axis of a bounded set Ω, which can be seen as an extension of the notion of Voronoi diagram to infinite sets. Interestingly, it is possible to construct certified approximations of the medial axis of quite general sets efficiently. One approach to be described consists in sampling the boundary of Ω and then computing an appropriate subset of the Voronoi diagram of the sample which approximates the medial

axis. Hence the problem of approximating the medial axis of Ω boils down to sampling the boundary of Ω, a problem that is closely related to mesh generation (see Chap. 5). Other informations on the medial axis can be found in Chap. 6.

2.7.1 Medial Axis and Lower Envelope

The *medial axis* of an open set Ω, denoted by $\mathcal{M}(\Omega)$, is defined as the the set of points of Ω that have more than one nearest neighbor on the boundary of Ω. Nearest refers in this section to the Euclidean distance although the results may be extended to other distance functions. A *medial sphere* σ is a sphere centered at a point c of the medial axis and passing through the nearest neighbors of c on $\partial\Omega$. Those points where σ is tangent (in the sense that σ does not enclose any point of $\partial\Omega$) to $\partial\Omega$ are called the *contact points* of σ.

The concept of medial axis can be considered as an extension of the notion of Voronoi diagram to infinite sets. Let o be a point of the boundary of Ω and δ_o be the distance function to o defined over Ω

$$\forall x \in \Omega : \ \delta_o(x) = \|x - o\|.$$

The *lower envelope* of the infinite set of functions δ_o is defined as

$$\Delta^- = \inf_{o \in \partial\Omega} \delta_o.$$

Following what we did for Voronoi diagrams (Sect. 2.2), we define, for any point $x \in \Omega$, its index set $I(x)$ as the set of all o such that $\Delta^-(x) = \delta_o(x)$. The set of points x such that $|I(x)| > 1$ constitutes the medial axis of Ω.

Computing the medial axis is difficult in general. If Ω is defined as a semi-algebraic set, i.e. a finite collection of algebraic equations and inequalities, $\mathcal{M}(\Omega)$ is also a semi-algebraic set that can therefore be computed using techniques from real algebraic geometry [32, 44]. This general approach, however, leads to algorithms of very high complexity. Theorem 1 can also be used but, still, working out the algebraic issues is a formidable task. Effective implementations are currently limited to simple objects. If Ω is a planar domain bounded by line segments and circular arcs, one can apply the results of Sect. 2.6. Further results can be found in [159].

An alternative and more practical approach consists in departing from the requirement to compute the medial axis exactly. In Sect. 2.7.2, we describe a method that approximates the medial axis of an object by first sampling its boundary, and then computing and pruning the Voronoi diagram of the sample.

2.7.2 Approximation of the Medial Axis

Approximating the medial axis of a set is a non trivial issue since sets that are close for the Hausdorff distance may have very different medial axes. This is

illustrated in Fig. 2.13. Let S be a closed curve, $\Omega = \mathbb{R}^2 \setminus S$ and $\mathcal{P}$ be a finite set of points approximating S. As can be seen on the figure, the skeleton of the Voronoi diagram of $\mathcal{P}$ is far from the medial axis of Ω since there are many long branches with no counterpart in $\mathcal{M}(\Omega)$. These branches are Voronoi edges whose dual Delaunay edges are small (their lengths tend to 0 when the sampling density increases). In other words, the medial axis is not continuous under the Hausdorff distance. Notice however that if we remove the long branches in Fig. 2.13, we obtain a good approximation of the medial axis of S. This observation will be made precise in Lemma 10 below. It leads to an approximate algorithm that first sample S and extract from the Euclidean Voronoi diagram of the sample a sub-complex that approximates the medial axis of S.

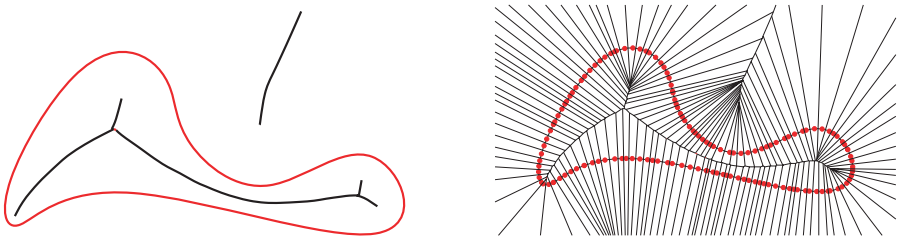

Fig. 2.13. On the left side, a closed curve S and the medial axis of $\Omega = \mathbb{R}^2 \setminus S$. On the right side, a dense sample E of points and its Voronoi diagram. The medial axis of $\mathbb{R}^2 \setminus E$ (which is the 1-skeleton of $\mathrm{Vor}(E)$) is very different from the medial axis of Ω

Given an open set Ω and a point x on the medial axis of Ω, we define $D(x)$ as the diameter of the smallest closed ball containing the contact points of the medial sphere centered at x. We define the λ-medial axis of Ω, denoted $\mathcal{M}_\lambda(\Omega)$ as the subset of the medial axis of Ω consisting of points x such that $D(x) \geq \lambda$.

Let Ω and Ω' be two open sets and let S and S' denote their boundaries. We assume that S and S' are compact and that their Hausdorff distance $d_H(S, S')$ is at most ε: any point of S is at distance at most ε from a point of S' and vice versa. Notice that we do not specify S nor S' to be finite or infinite point sets. For convenience, we rename medial axis of S (resp. S') and write $\mathcal{M}(S)$ (resp. $\mathcal{M}(S')$) the medial axis of $\mathbb{R}^3 \setminus S$. Similarly, we rename medial axis of S' and write $\mathcal{M}(S')$) the medial axis of $\mathbb{R}^3 \setminus S$.

The following lemma says that the λ-medial axis of S is close to the medial axis of S', provided that λ is sufficiently large. It should be emphasized that close here refers to the one-sided Hausdorff distance: the medial axis of S' is not necessarily close to the λ-medial axis of S although, by exchanging the roles of S and S', the lemma states that the λ'-medial axis of S' is close to

the medial axis of S for a sufficiently large λ'. We will go back to this point later.

We say that a ball is S-empty if its interior does not intersect S. The sphere bounding a S-empty ball is called a S-empty sphere.

Lemma 10. *Let σ be a S-empty sphere centered at c, of radius r, intersecting S in two points x and y. If $\varepsilon < \frac{r}{2}$ and $l \stackrel{\text{def}}{=} \frac{\|x-y\|}{4} \geq \sqrt{\varepsilon r(1-\frac{\varepsilon}{r})}$, there exists an S'-empty sphere tangent to S' in two points whose center c' and radius r' are such that $|1-\frac{r'}{r}|$ and $\frac{\|c'-c\|}{r}$ are at most $\delta = \frac{\varepsilon r}{l^2-\varepsilon r+\varepsilon^2}$.*

Proof. Let σ'' be the maximal S'-empty sphere centered at c and let r'' be its radius (see Fig. 2.14). The Hausdorff distance between S and S' being at most ε, we have $|r-r''| \leq \varepsilon$. Let y' be a point of $\sigma'' \cap S'$.

Let σ' be the maximal S'-empty sphere tangent to σ'' at y'. σ' is tangent to S' at at least two points. Let c' be its center and r' its radius.

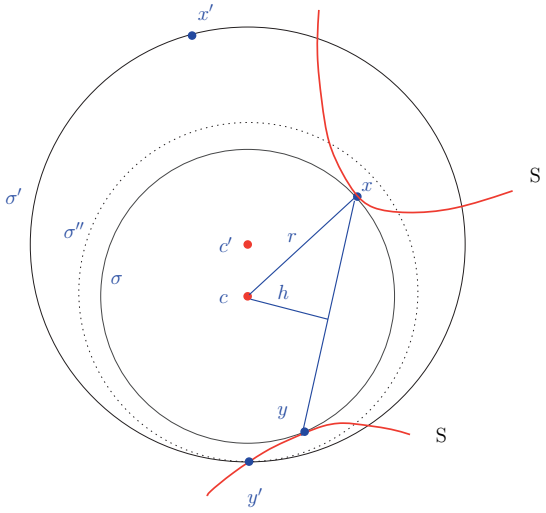

Fig. 2.14. S is the continuous curve. x' and y' belong to S'

Noting $h = \sqrt{r^2 - \frac{1}{4}\|x-y\|^2} = \sqrt{r^2-4l^2}$ the distance from c to line xy, we have

$$\begin{aligned} d(c', S) &\leq \min(\|c'-x\|, \|c'-y\|) \\ &\leq \sqrt{(\|c-c'\|+h)^2 + \frac{1}{4}\|x-y\|^2} \\ &= \sqrt{r^2 + \|c-c'\|^2 + 2h\|c-c'\|}. \end{aligned}$$

On the other hand,

$$d(c', S') = \|c' - y'\| = \|c - c'\| + r'' \geq \|c - c'\| + r - \varepsilon.$$

From these two inequalities, we deduce

$$\|c - c'\| + r - \varepsilon \leq d(c', S') \leq d(c', S) + \varepsilon \leq \sqrt{r^2 + \|c - c'\|^2 + 2h\|c - c'\|} + \varepsilon.$$

and, since, by assumption, $r > 2\varepsilon$, we get $\|c - c'\|\,(r - 2\varepsilon - h) \leq 2\varepsilon(r - \varepsilon)$. Moreover, by assumption, $l \geq \sqrt{\varepsilon r(1 - \frac{\varepsilon}{r})}$ which implies $r - 2\varepsilon - h \geq 0$. Indeed,

$$h = \sqrt{r^2 - 4l^2} \leq \sqrt{r^2 - 4\varepsilon r + 4\varepsilon^2} = r - 2\varepsilon.$$

We then deduce

$$\|c - c'\| \leq \frac{2\varepsilon(r - \varepsilon)}{r - h - 2\varepsilon} \leq \frac{2\varepsilon r}{r - \sqrt{r^2 - 4l^2} - 2\varepsilon}.$$

We then get

$$\frac{\|c - c'\|}{r} \leq \frac{2\varepsilon(r + \sqrt{r^2 - 4l^2} - 2\varepsilon)}{(r - 2\varepsilon)^2 - (r^2 - 4l^2)} \leq \frac{\varepsilon r}{l^2 - \varepsilon r + \varepsilon^2}.$$

The same bound plainly holds for $|1 - \frac{r'}{r}|$.

We consider the case where S is a surface of $\mathbb{R}^3$ and S' is a finite set of points on S at Hausdorff distance at most ε from S. To avoid confusion, we rename S' as $\mathcal{P}$. As already noticed, $\mathcal{M}(\mathcal{P})$ is the 1-skeleton of the Voronoi diagram $\mathrm{Vor}(\mathcal{P})$, called simply the Voronoi diagram of $\mathcal{P}$ in this section. The set of contact points of a medial sphere centered at a point $c \in \mathcal{M}(\mathcal{P})$ is the set of closest points of c in $\mathcal{P}$. Any point in the relative interior of a face of $\mathrm{Vor}(\mathcal{P})$ has the same closest points in $\mathcal{P}$. It follows that the λ-medial axis of $\mathcal{P}$ is the subset of the faces of $\mathrm{Vor}(\mathcal{P})$ whose contact points cannot be enclosed in a ball of diameter λ.

Lemma 10 then says that any Delaunay sphere of $\mathrm{Del}(\mathcal{P})$ passing through two sample points that are sufficiently far apart, is close to a medial sphere of S (for the Hausdorff distance). We have therefore bounded the *one-sided* Hausdorff distance from the λ-medial axis (Voronoi diagram) of an ε-sample $\mathcal{P}$ of S to the medial axis of S, when λ is sufficiently large with respect to ε. If we apply the lemma the other way around, we see that, for sufficiently large λ', $\mathcal{M}_{\lambda'}(S)$ is close to $\mathcal{M}(\mathcal{P})$. However, as observed above (Fig. 2.13), we cannot hope to bound the two-sided Hausdorff distance between $\mathcal{M}(S)$ and $\mathcal{M}(\mathcal{P})$.

The above lemma can be strengthened as recently shown by Chazal and Lieutier [83, 32]. They proved that the λ-medial axis of S is close to the λ'-medial axis of $\mathcal{P}$ for a sufficiently large λ and some positive λ' that depends on λ and ε. More precisely, let D be the diameter of S and k'' a positive constant. They showed that there exist three functions of ε, $k(\varepsilon) = 15\sqrt{2}\,\sqrt[4]{D^3\varepsilon}$, $k'(\varepsilon) = 10\sqrt{3}\,\sqrt[4]{D^3\varepsilon}$ and $k''(\varepsilon) = k''\sqrt[4]{D^3\varepsilon}$, such that

$$\mathcal{M}_{k(\varepsilon)}(S) \subset \mathcal{M}_{k'(\varepsilon)}(\mathcal{P}) \oplus B_{2\sqrt{D\varepsilon}} \subset \mathcal{M}_{k''(\varepsilon)}(S) \oplus B_{4\sqrt{D\varepsilon}}.$$

Here, B_r denotes the ball centered at the origin of radius r, and $\oplus$ the Minkowski sum.

Consider now a family of point sets $\mathcal{P}_\varepsilon$ parametrized by ε such that $d_H(S, \mathcal{P}_\varepsilon) \leq \varepsilon$ and let ε tends to 0. Because $\mathcal{M}_\eta(S)$ tends to $\mathcal{M}(S)$ when η tends to 0, we deduce from the above inequalities, that

$$\lim_{\varepsilon \to 0} d_H(\mathcal{M}(S), \mathcal{M}_{10\sqrt[4]{9D^3\varepsilon}}(\mathcal{P}_\varepsilon)) = 0.$$

The above discussion provides an algorithm to approximate the medial axis of S within any specified error (see Algorithm 5).

Algorithm 5 Approximation of the Medial Axis

INPUT: A surface S and a positive real ε

1. Sample S so as to obtain a sample $\mathcal{P}$ such that $d_H(S, \mathcal{P}) \leq \varepsilon$;
2. Construct the Voronoi diagram of $\mathcal{P}$;
3. Remove from the diagram the faces for which the diameter of the set of contact points is smaller than $10\sqrt[4]{9D^3\varepsilon}$.

OUTPUT: A PL approximation of $\mathcal{M}(S)$

The main issue is therefore to compute a sample of points on S (step 1). If S is a surface of $\mathbb{R}^3$, one can use a surface mesh generator to mesh S and take for $\mathcal{P}$ the vertices of the mesh. Various algorithms can be found in Chap. 5 and we refer to that chapter for a thorough description and analysis of these algorithms. Especially attractive in the context of medial axis approximation, are the algorithms that are based on the 3-dimensional Delaunay triangulation since we get the Voronoi diagram of the sample points (step 2) at no additional cost. An example obtained with the surface mesh generator of Boissonnat and Oudot [65] is shown in Fig. 2.15.

Exercise 31. Let O be a bounded open set. Show that $\mathcal{M}(O)$ is a retract of O (and therefore has the same homotopy type as O) [243].

2.8 Voronoi Diagrams in CGAL

The Computational Geometry Algorithms Library CGAL [2] offers severals packages to compute Voronoi diagrams. Euclidean Voronoi diagrams of points and power diagrams are represented through their dual Delaunay and regular triangulations. CGAL provides Delaunay and regular triangulations in $\mathbb{R}^2$ and $\mathbb{R}^3$. The implementation is based on a randomized incremental algorithm

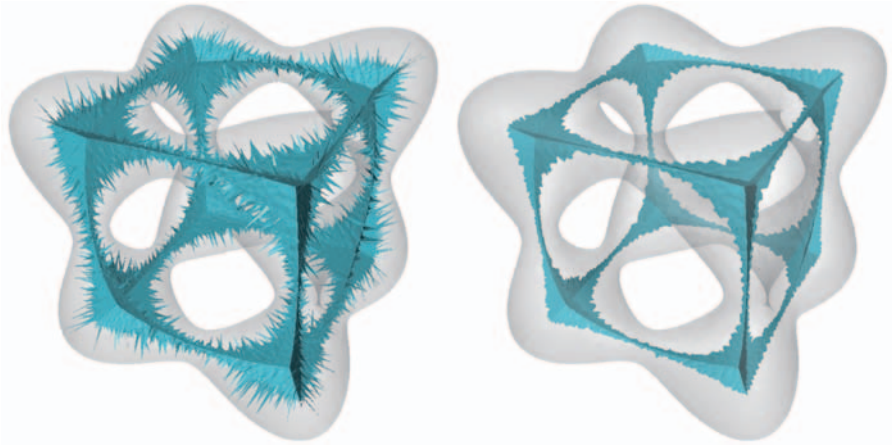

Fig. 2.15. Two λ-medial axes of the same shape, with λ increasing from left to right, computed as a subset of the Voronoi diagram of a sample of the boundary (courtesy of Steve Oudot)

using a variant of the Voronoi hierarchy described in Sect. 2.6. Delaunay triangulations are also provided in higher dimensions.

The library also contains packages to compute Voronoi diagrams of line segments [215] and Apollonius diagrams in $\mathbb{R}^2$ [216]. Those packages implement the incremental algorithm described in Sect. 2.6. A prototype implementation of Möbius diagrams in $\mathbb{R}^2$ also exists. This prototype computes the Möbius diagram as the projection of the intersection of a 3-dimensional power diagram with a paraboloid, as described in Sect. 2.4.1. This prototype also serves as the basis for the developement of a CGAL package for 3-dimensional Apollonius diagrams, where the boundary of each cell is computed as a 2-dimensional Möbius diagram, following the results of Sect. 2.4.3 [62]. See Fig. 2.8.

2.9 Applications

Euclidean and affine Voronoi diagrams have numerous applications we do not discuss here. The interested reader can consult other chapters of the book, most notably Chap. 5 on surface meshing and Chap. 6 on reconstruction. Other applications can be found in the surveys and the textbooks mentionned in the introduction.

Additively and multiplicatively weighted distances arise when modeling growing processes and have important applications in biology, ecology and other fields. Consider a number of crystals, all growing at the same rate, and all starting at the same time : one gets a number of growing circles. As these circles meet, they draw a Euclidean Voronoi diagram. In reality, crystals start

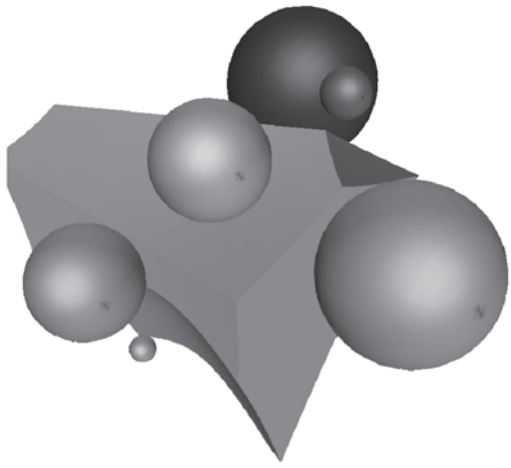

Fig. 2.16. A cell in an Apollonius diagram of spheres

growing at different times. If they still grow at the same rate, they will meet along an Apollonius diagram. This growth model is known as the Johnson-Mehl model in cell biology. In other contexts, all the crystals start at the same time, but grow at different rates. Now we get what is called the multiplicatively weighted Voronoi diagram, a special case of Möbius diagrams.

Spheres are common models for a variety of objects such as particles, atoms or beads. Hence, Apollonius diagrams have been used in physics, material sciences, molecular biology and chemistry [245, 339, 227, 228]. They have also been used for sphere packing [246] and shortest paths computations [256].

Euclidean Voronoi diagrams of non punctual objects find applications in robot motion planning [237, 197]. Medial axes are used for shape analysis [160], for computing offsets in Computer-Aided Design [118], and for mesh generation [290, 289, 316]. Medial axes are also used in character recognition, road network detection in geographic information systems, and other applications.

Acknowledgments

We thank D. Attali, C. Delage and M. Karavelas with whom part of the research reported in this chapter has been conducted. We also thank F. Chazal and A. Lieutier for fruitful discussions on the approximation of the medial axis.