Chapter 17

Euclidean metric

This chapter is concerned with the simplest case of Voronoi diagrams, where the objects are points and the distance is given by the usual Euclidean metric in $\mathbb{E}^d$. The cells in the Voronoi diagram of a set $\mathcal{M}$ of points are then the equivalence classes of the equivalence relation "to have the same nearest neighbor in $\mathcal{M}$". It is possible to show (see section 17.2) that such cells can be obtained by projecting the facets of a polytope in $\mathbb{E}^{d+1}$ onto $\mathbb{E}^d$, which enables us to use several results concerning polytopes for Voronoi diagrams as well. Bounds can be obtained in this way for the complexity of Voronoi diagrams and of their computation. In section 17.3, we define a dual of the Voronoi diagram, the Delaunay complex, that enjoys several properties which make it desirable in applications such as numerical analysis in connection with finite-element methods. The last section of this chapter introduces a first generalization of Voronoi diagrams (see section 17.4): the higher-order Voronoi diagrams. The cells in the diagram of order k are the equivalence classes of the equivalence relation "to have the same k nearest neighbors in $\mathcal{M}$", a notion that is often very helpful in data analysis.

17.1 Definition

Let $\mathcal{M}$ be a set of n points in $\mathbb{E}^d$, $M_1, \ldots, M_n$, which we call the *sites* to avoid confusion with the other points in $\mathbb{E}^d$. To each site M_i we attach the region $V(M_i)$ in $\mathbb{E}^d$ that contains the points in $\mathbb{E}^d$ closer to M_i than to any other point in $\mathcal{M}$:

$$V(M_i) = \{X \in \mathbb{E}^d \ : \ \delta(X, M_i) \leq \delta(X, M_j) \text{ for any } j \neq i\}.$$

In this chapter, δ denotes the Euclidean distance in $\mathbb{E}^d$. Other distances will be considered in chapter 18.

The set of points closer to M_i than to another site M_j is the half-space that contains M_i and that is bounded by the *perpendicular bisector* of the segment

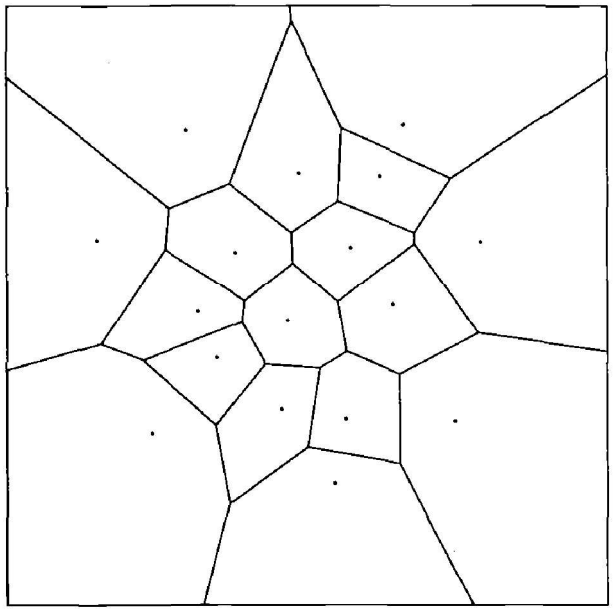

Figure 17.1. The Voronoi diagram of a set of points in the plane.

M_iM_j: this is the hyperplane perpendicular to M_iM_j that intersects M_iM_j at the midpoint of M_i and M_j. The region $V(M_i)$ is thus the intersection of a finite number of closed half-spaces, bounded by the perpendicular bisectors of M_iM_j, $j = 1, \ldots, n$, $j \neq i$. This shows that $V(M_i)$ is a convex polytope, which may or may not be bounded. As we will see later, the $V(M_i)$'s and their faces form a cell complex whose domain is the whole of $\mathbb{E}^d$. This complex is called the *Voronoi diagram* of $\mathcal{M}$ and is denoted by $\mathcal{V}or(\mathcal{M})$ (see figure 17.1).

A first and useful interpretation of the Voronoi diagram (another interpretation is given in the next chapter) views the cell $V(M_i)$ as the set of centers of balls such that the boundary of such a ball contains M_i and its interior does not contain another site M_j, $j \neq i$. In particular, this point of view leads to the interpretation of a Voronoi diagram in $\mathbb{E}^d$ as a polytope in $\mathbb{E}^{d+1}$, which also enables it to be computed efficiently. This interpretation is developed in section 17.2 where we represent spheres of $\mathbb{E}^d$ as points in $\mathbb{E}^{d+1}$.

From now on, we say that the sites are in *L_2-general* position if no sphere can contain $d + 2$ sites on its boundary.

17.2 Voronoi diagrams and polytopes

17.2.1 Power of a point with respect to a sphere

Consider the Euclidean space of dimension d, $\mathbb{E}^d$, and let O be its origin, and Σ be a sphere of $\mathbb{E}^d$ centered at C with radius r. Its equation is given by $\Sigma(X) = 0$

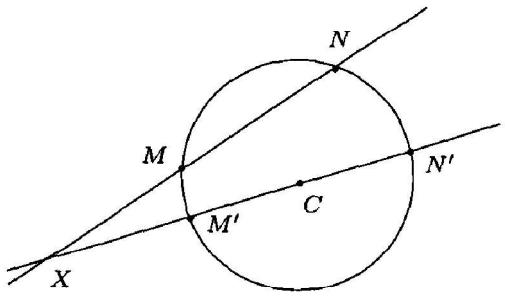

Figure 17.2. Power of a point with respect to a circle.

where

$$\Sigma(X) = XC^2 - r^2. \tag{17.1}$$

By the *interior* of a sphere Σ, we mean the set of points X such that $\Sigma(X)$ is negative. The *exterior* is the set of points X such that $\Sigma(X)$ is negative. A point X is said to be *on*, *inside* or *outside* a sphere if it belongs to the sphere, respectively to its interior, to its exterior. For any point X in $\mathbb{E}^d$, $\Sigma(X)$ is called the *power* of X with respect to Σ. The power of the origin with respect to Σ is also denoted by σ and we have

$$\sigma = \Sigma(O) = C^2 - r^2. \tag{17.2}$$

If D is any line that contains X, and if M and N are the intersection points of D with Σ, then

$$\Sigma(X) = XM \cdot XN. \tag{17.3}$$

This is obvious when D is the line connecting X and C. Otherwise let D' be the line that contains X and C, and let M' and N' be its intersection points with Σ (see figure 17.2). The triangles XMM' and $XN'N$ are similar (the angles $\widehat{M'MN}$ and $\widehat{M'N'N}$ are supplementary), which proves equation 17.3. In the case where X belongs to the exterior of Σ and D is tangent to Σ at T, then $M = N = T$ and the previous equation can be rewritten

$$\Sigma(X) = XT^2. \tag{17.4}$$

17.2.2 Representation of spheres

Let ϕ be the mapping that takes a sphere Σ in $\mathbb{E}^d$, of center C and whose power with respect to O is σ, to the point $\phi(\Sigma) = (C, \sigma)$ in $\mathbb{E}^{d+1}$. Using ϕ enables us to treat spheres in $\mathbb{E}^d$ just as points in $\mathbb{E}^{d+1}$.

We embed $\mathbb{E}^d$ as the hyperplane in $\mathbb{E}^{d+1}$ whose equation is $x_{d+1} = 0$. As usual, the direction of the x_{d+1}-axis is called the vertical direction and we use the words

above and *below* in connection with this vertical ordering. We denote by X a point of $\mathbb{E}^d$ or its coordinate vector $(x_1, \ldots, x_d)$ indifferently, and by $\underline{X}$ a point in $\mathbb{E}^{d+1}$ or its coordinate vector $(x_1, \ldots, x_{d+1})$. By the above embedding, $\phi(\Sigma)$ projects vertically onto C. Later on, we will often use vertical projections and, unless mentioned otherwise, the word *projection* refers to the vertical projection from $\mathbb{E}^{d+1}$ onto $\mathbb{E}^d$.

We also use homogeneous coordinates and the matrix notation. We denote by $\boldsymbol{X} = (x_1, \ldots, x_d, t)$ (resp. $\underline{\boldsymbol{X}} = (x_1, \ldots, x_{d+1}, t)$) the homogeneous coordinate vector of a point X in $\mathbb{E}^d$ (resp. a point $\underline{X}$ in $\mathbb{E}^{d+1}$). The equation of the sphere Σ can then be rewritten with homogeneous coordinates as

$$\boldsymbol{X}\,\boldsymbol{\Sigma}\,\boldsymbol{X}^t = 0 \quad \text{with} \quad \boldsymbol{\Sigma} = \begin{pmatrix} \mathbb{I}_d & -C^t \\ -C & \sigma \end{pmatrix},$$

where $\mathbb{I}_d$ denotes the $d \times d$ identity matrix.

17.2.3 The paraboloid $\mathcal{P}$

From equation 17.2, it follows that the images under ϕ of points in $\mathbb{E}^d$, considered as spheres of radius 0, belong to the paraboloid of revolution $\mathcal{P}$ with vertical axis and equation

$$x_{d+1} = \sum_{i=1}^{d} x_i^2 = X \cdot X \quad \text{with} \quad X = (x_1, \ldots, x_d).$$

In a homogeneous system of coordinates $\mathcal{P}$ is given by

$$\underline{\boldsymbol{X}}\,\Delta_{\mathcal{P}}\,\underline{\boldsymbol{X}}^t = 0 \quad \text{where} \quad \Delta_{\mathcal{P}} = \begin{pmatrix} \mathbb{I}_d & 0 & 0 \\ 0 & 0 & -1/2 \\ 0 & -1/2 & 0 \end{pmatrix}.$$

Identifying a point X and the sphere centered at X with radius 0 shows that ϕ maps any point X in $\mathbb{E}^d$ to the point $\phi(X)$ in $\mathbb{E}^{d+1}$ obtained by lifting X onto $\mathcal{P}$.

The set of concentric spheres in $\mathbb{E}^d$, centered at C, is mapped by ϕ onto the vertical line in $\mathbb{E}^{d+1}$ that contains C (and hence $\phi(C)$). Let Σ be such a sphere. Equation 17.2 implies that the signed vertical distance from $\phi(\Sigma)$ to $\phi(C)$ equals r^2 (see figure 17.3). Thus, the *real* spheres, whose squared radii are non-negative, are mapped by ϕ to the points lying on or below the paraboloid, while the points lying above the paraboloid are the images under ϕ of the *imaginary* spheres, whose squared radii are negative.

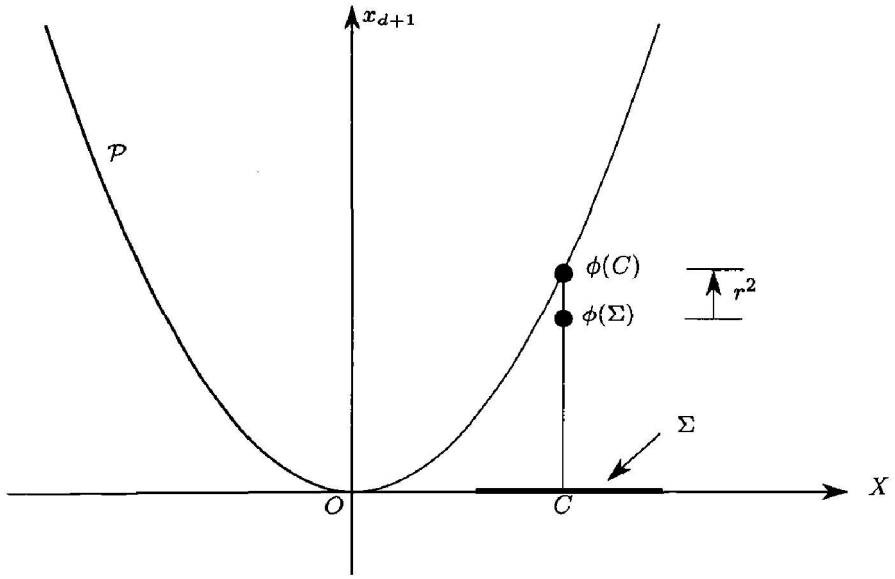

Figure 17.3. The paraboloid $\mathcal{P}$ (the coordinate system is not normed, so as to simplify the representation).

17.2.4 Polarity

Consider a quadric $\mathcal{Q}$ in $\mathbb{E}^{d+1}$ defined by its homogeneous equation

$$\mathcal{Q}(\underline{X}) = \underline{X}\Delta_{\mathcal{Q}}\underline{X}^t = 0.$$

Henceforth, $\mathcal{Q}$ will be the paraboloid $\mathcal{P}$, but we treat the case of any quadric for generality as it introduces no additional difficulty. Two points X and Y are *conjugate* with respect to $\mathcal{Q}$ if

$$\mathcal{Q}(\underline{X}, \underline{Y}) = \underline{X}\Delta_{\mathcal{Q}}\underline{Y}^t = 0.$$

The polarity with respect to $\mathcal{Q}$ described in section 7.3 is an involution between points and hyperplanes in $\mathbb{E}^{d+1}$ which maps any point $\bar{A}$ to its *polar hyperplane* $\bar{A}^*$ of equation

$$\underline{A}\Delta_{\mathcal{Q}}\underline{X}^t = 0,$$

and maps any hyperplane $\underline{H}$ to a point $\underline{H}^*$ whose polar hyperplane is $\underline{H}$. The point $\underline{H}^*$ is called the *pole* of $\underline{H}$.

Note that if $\mathcal{Q}$ is the paraboloid $\mathcal{P}$ and if we put $\phi(\Sigma) = (C, \sigma)$, the equation of the polar hyperplane $\phi(\Sigma)^*$ of $\phi(\Sigma)$ can be rewritten as

$$x_{d+1} = 2C \cdot X - \sigma.$$

An essential property of polarity is that it preserves incidences (see section 7.3): a point $\underline{X}$ belongs to a hyperplane $\underline{H}$ if and only if its polar hyperplane $\underline{X}^*$

contains the pole $\underline{H}^*$ of $\underline{H}$. Moreover (see exercise 7.14), when the quadric is the paraboloid $\mathcal{P}$, we have

$$\begin{aligned} \underline{X} \in \underline{H}^+ &\iff \underline{H}^* \in \underline{X}^{*+} \\ \underline{X} \in \underline{H}^- &\iff \underline{H}^* \in \underline{X}^{*-}, \end{aligned}$$

if we denote by $\underline{H}^+$ and $\underline{H}^-$ the half-spaces bounded by $\underline{H}$ and that lie respectively above and below $\underline{H}$.

17.2.5 Orthogonal spheres

Two spheres Σ_1 and Σ_2 centered at C_1 and C_2 and with radii r_1 and r_2 are *orthogonal* if

$$\Sigma_1(C_2) = r_2^2, \tag{17.5}$$

or equivalently if

$$\Sigma_2(C_1) = r_1^2.$$

A simple verification shows that, if the spheres are real, then they are orthogonal if and only if the angle (IC_1, IC_2) at any intersection point I of $\Sigma_1 \cap \Sigma_2$ is a right angle, or equivalently, if and only if the dihedral angle of the tangent hyperplanes at I is a right angle.

Expression 17.5 may be rewritten as

$$C_1 \cdot C_2 - \frac{1}{2}(\sigma_1 + \sigma_2) = 0, \quad \text{with } \sigma_i = C_i^2 - r_i^2 \ \ (i = 1, 2),$$

which shows that two spheres Σ_1 and Σ_2 are orthogonal if the two points $\phi(\Sigma_1)$ and $\phi(\Sigma_2)$ are conjugate with respect to the paraboloid $\mathcal{P}$. This implies that:

Lemma 17.2.1 *The set of spheres in $\mathbb{E}^d$ that are orthogonal to a given sphere is mapped by ϕ to the polar hyperplane $\phi(\Sigma)^*$ of $\phi(\Sigma)$.*

Let us now consider the points in $\mathbb{E}^d$ as spheres of radius 0. The set of spheres in $\mathbb{E}^d$ that pass through a given point $X \in \mathbb{E}^d$ is also the set of spheres orthogonal to the sphere centered at X with radius 0. Therefore its image under ϕ is the hyperplane $\phi(X)^*$ polar to $\phi(X) \in \mathcal{P}$. This hyperplane must be tangent to $\mathcal{P}$ and to $\phi(X)$: indeed, the only sphere of radius 0 which is orthogonal to X is X itself, and hence $\phi(X)^*$ intersects $\mathcal{P}$ in a single point $\phi(X)$.

Let Σ be a sphere in $\mathbb{E}^d$. The intersection of $\phi(\Sigma)^*$ with $\mathcal{P}$ is the image under ϕ of the set of spheres with radius 0 that are orthogonal to Σ, namely Σ itself (considered as a set of points, or equivalently as a set of spheres of radius 0). Consequently, $\phi(\Sigma)^* \cap \mathcal{P}$ in $\mathbb{E}^{d+1}$ projects onto Σ in $\mathbb{E}^d$. More generally, we have the following result.

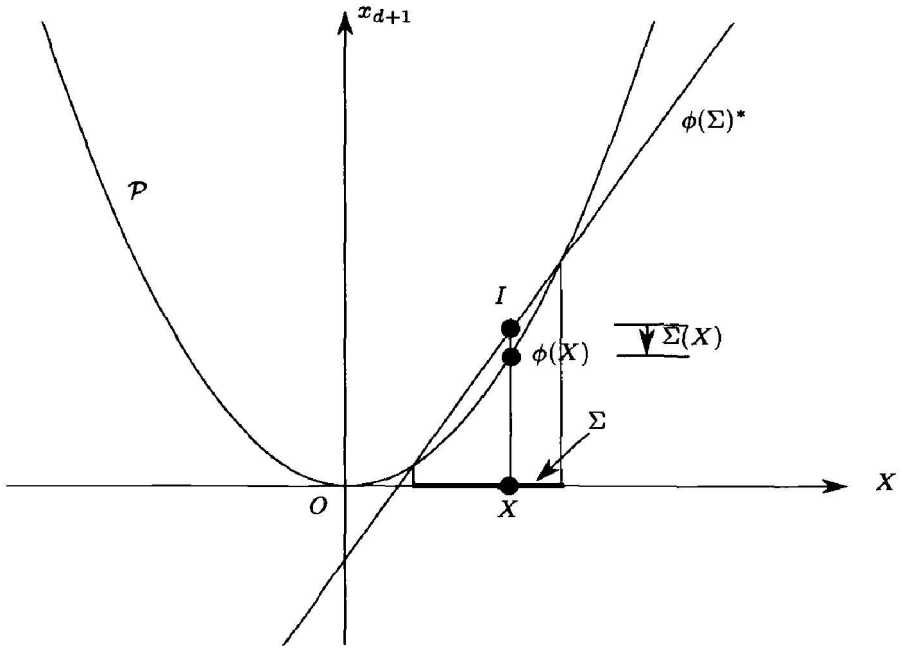

Figure 17.4. Interpretation in $\Sigma(X)$ in $\mathbb{E}^{d+1}$.

Lemma 17.2.2 *The intersection of the paraboloid $\mathcal{P}$ with a hyperplane $\underline{H}$ projects onto $\mathbb{E}^d$ as a sphere $\phi^{-1}(\underline{H}^*)$ whose center is the vertical projection of $\underline{H}^*$. Conversely, the points of a sphere Σ of $\mathbb{E}^d$ lifted on the paraboloid $\mathcal{P}$ in $\mathbb{E}^{d+1}$ belong to a unique hyperplane that intersects $\mathcal{P}$ exactly at these points. This hyperplane is the polar hyperplane $\phi(\Sigma)^*$ of $\phi(\Sigma)$.*

It follows from this lemma that the power of a point X with respect to a sphere Σ equals the square of the radius of the sphere Σ_X orthogonal to Σ and centered at X (Σ_X is imaginary if X is inside Σ). The power $\Sigma(X)$ can be easily computed in the space $\mathbb{E}^{d+1}$ that represents the spheres of $\mathbb{E}^d$. Indeed (see figure 17.4), Σ_X is mapped by ϕ to a point $\underline{I}$ in $\mathbb{E}^{d+1}$ that is the intersection of the vertical line that passes through X (which corresponds to the spheres centered at X) with the polar hyperplane $\phi(\Sigma)^*$ of $\phi(\Sigma)$ (which corresponds to the spheres orthogonal to Σ). The x_{d+1}-coordinates of $\phi(x)$ and $\underline{I}$ are respectively X^2 and $\Sigma_X(O) = X^2 - \Sigma(X)$ since the square of the radius of Σ_X equals the power of X with respect to Σ. The difference of these x_{d+1}-coordinates is called the *signed vertical distance.* This proves the following lemma.

Lemma 17.2.3 *The power of X with respect to a sphere Σ equals the signed vertical distance from the point $\phi(X)$ to the hyperplane $\phi(\Sigma)^*$.*

We thus have the following lemma:

Lemma 17.2.4 *Let X and Σ be respectively a point and a sphere in $\mathbb{E}^d$. If $\underline{H}$ is a hyperplane in $\mathbb{E}^{d+1}$, we denote by $\underline{H}^-$ the half-space lying below $\underline{H}$. Then:*

$$\begin{array}{rcccl} X \in \Sigma & \Longleftrightarrow & \phi(X) \in \phi(\Sigma)^* & \Longleftrightarrow & \phi(\Sigma) \in \phi(X)^* \\ X \in int(\Sigma) & \Longleftrightarrow & \phi(X) \in \phi(\Sigma)^{*-} & \Longleftrightarrow & \phi(\Sigma) \in \phi(X)^{*-} \\ X \in ext(\Sigma) & \Longleftrightarrow & \phi(X) \in \phi(\Sigma)^{*+} & \Longleftrightarrow & \phi(\Sigma) \in \phi(X)^{*+} \end{array}$$

The equivalences on the left are consequences of the two preceding lemmas, and the ones on the right are proved by the special properties of polarity (see subsection 17.2.4 and exercise 7.13).

Any point in the half-space that lies below $\phi(X)^*$ in $\mathbb{E}^{d+1}$ is thus the image under ϕ of a sphere whose interior contains X. Likewise, any point in the half-space that lies above $\phi(X)^*$ in $\mathbb{E}^{d+1}$ is the image under ϕ of a sphere whose exterior contains X, and the points on $\phi(X)^*$ are the images of the spheres passing through X.

Remark. Lemma 17.2.3 shows that the squared distance $\|XA\|^2$ separating points X and A, which is also the power of X with respect to the sphere centered at A with radius 0, equals the absolute value of the vertical distance between $\phi(A)^*$ and $\phi(X)$. Points X and A play symmetric roles, so $\|XA\|^2$ also equals the absolute value of the vertical distance between $\phi(X)^*$ and $\phi(A)$.

17.2.6 Radical hyperplane

Let Σ_1 and Σ_2 be two spheres in $\mathbb{E}^d$. The set of points in $\mathbb{E}^d$ that have the same power with respect to these two spheres is a hyperplane, called the *radical hyperplane* and denoted by H_{12}, whose equation is given by

$$H_{12}: \quad \Sigma_1(X) - \Sigma_2(X) = 0.$$

As we observed in subsection 17.2.5, the power of a point X with respect to a sphere Σ equals the square of the radius of the sphere orthogonal to Σ centered at X. A point has the same power with respect to Σ_1 as with respect to Σ_2 if it is the center of a sphere orthogonal to both Σ_1 and Σ_2. Lemma 17.2.1 shows that the set of spheres in $\mathbb{E}^d$ that are orthogonal to a given sphere Σ is mapped by ϕ onto the polar hyperplane $\phi(\Sigma)^*$. The spheres orthogonal to Σ_1 and Σ_2 are thus mapped by ϕ to the affine subspace of dimension $d-1$ that is the intersection of $\phi(\Sigma_1)^*$ and $\phi(\Sigma_2)^*$. The projection onto $\mathbb{E}^d$ of this affine subspace is exactly the set of points that have the same power with respect to Σ_1 and Σ_2.

17.2.7 Voronoi diagrams

Let $\mathcal{M} = \{M_1, \ldots, M_n\}$ be a set of n points in $\mathbb{E}^d$. As before, we embed the Euclidean space $\mathbb{E}^d$ of dimension d into $\mathbb{E}^{d+1}$ as the hyperplane $x_{d+1} = 0$, and

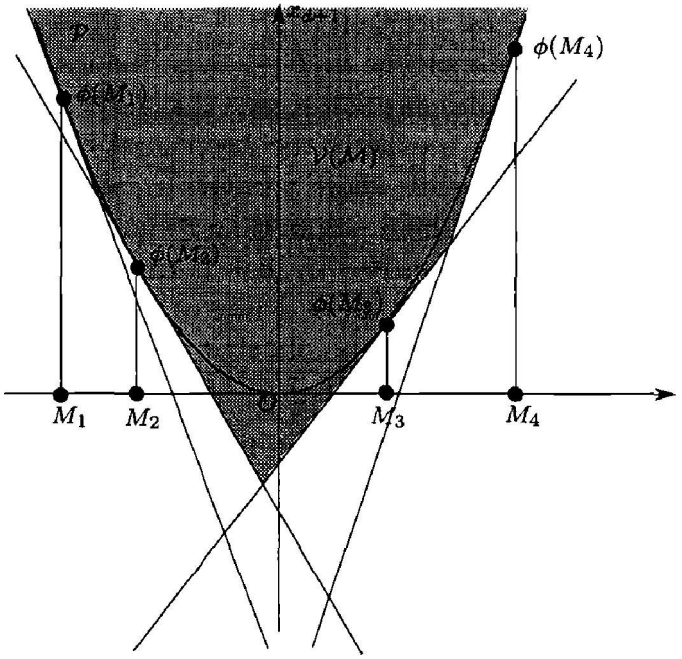

Figure 17.5. The Voronoi polytope $\mathcal{V}(\mathcal{M})$.

we let $\phi(M_i)^*$ denote the hyperplane in $\mathbb{E}^{d+1}$ that is tangent to the paraboloid $\mathcal{P}$ at the point $\phi(M_i)$ obtained by lifting M vertically onto the paraboloid $\mathcal{P}$, for each $i = 1, \ldots, n$. The preceding discussion shows that the set of spheres (real or imaginary) whose interiors contain no point of $\mathcal{M}$ is mapped by ϕ to the intersection of the n half-spaces lying above the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$. This intersection is an unbounded polytope which contains $\mathcal{P}$. We call it the *Voronoi polytope* and denote it by $\mathcal{V}(\mathcal{M})$ (see figure 17.5).

Theorem 17.2.5 *The Voronoi diagram of $\mathcal{M}$, denoted by $\mathcal{V}or(\mathcal{M})$, is a cell complex of dimension d in $\mathbb{E}^d$ whose faces are obtained by projecting onto $\mathbb{E}^d$ the proper faces of the Voronoi polytope $\mathcal{V}(\mathcal{M})$.*

Proof. The boundary of $\mathcal{V}(\mathcal{M})$ is a pure cell complex of dimension d, hence so is $\mathcal{V}or(\mathcal{M})$. Let $\overline{A}$ be a point on a facet of $\mathcal{V}(\mathcal{M})$ that is contained in the hyperplane tangent to $\mathcal{P}$ at $\phi(M_i)$. Then $\overline{A}$ is the image under ϕ of a sphere Σ_A that passes through M_i and whose interior contains no other point of $\mathcal{M}$ (see lemma 17.2.4). There cannot be a site in $\mathcal{M}$ closer to the center of Σ_A than M_i. But this center is exactly the projection A of $\overline{A}$ onto $\mathbb{E}^d$. In other words, A belongs to the cell $V(M_i)$ of the Voronoi diagram. □

This theorem implies that the combinatorial properties of Voronoi diagrams follow directly from those of polytopes as studied in chapter 7. In particular, if

the points $\phi(M_i)^*$ are in general position in $\mathbb{E}^{d+1}$, then $\mathcal{V}(\mathcal{M})$ is a simple $(d+1)$-polytope. Each vertex is thus incident to $d+1$ hyperplanes. Expressed in terms of M_i's, the *general condition assumption* means that no $d+2$ points in $\mathcal{M}$ lie on the boundary of a sphere: this is exactly the L_2-general position assumption. If it is satisfied, $\mathcal{V}or(\mathcal{M})$ is a complex whose vertices are all equidistant from some $d+1$ points in $\mathcal{M}$ and closer to these points than to any other point in $\mathcal{M}$: they are the centers of spheres circumscribed to $(d+1)$-tuples whose interiors do not contain any point in $\mathcal{M}$. More generally, a k-face of $\mathcal{V}or(\mathcal{M})$ is the projection of a k-face of $\mathcal{V}(\mathcal{M})$. It is thus the set of points that are equidistant from $d+1-k$ points in $\mathcal{M}$ and closer to these points than to any other point in $\mathcal{M}$.

Theorem 17.2.5 reduces the problem of computing the Voronoi diagram of n points in $\mathbb{E}^d$ to the computation of the intersection of n half-spaces of $\mathbb{E}^d$. The algorithms described in this book that compute half-space intersections, be they deterministic, randomized, static or dynamic, output-sensitive or not, can all be used to compute Voronoi diagrams.

Corollary 17.2.6 *The complexity (namely, the number of faces) of the Voronoi diagrams of n points in $\mathbb{E}^d$ is $\Theta(n^{\lceil d/2 \rceil})$. We may compute such a diagram in time $O(n \log n + n^{\lceil d/2 \rceil})$, which is optimal in the worst case.*

Proof. The upper bounds on the complexity and running time of the algorithm are immediate consequences of the upper bound theorem 7.2.5 and of results of the previous sections.

That $\Omega(n^{\lceil d/2 \rceil})$ is a lower bound on the complexity of the Voronoi diagram of n points in $\mathbb{E}^d$ is a consequence of exercise 7.11, where it is shown how to construct a maximal polytope whose vertices lie on the paraboloid, and of theorem 17.3.1 below.

That $\Omega(n \log n)$ is a lower bound on computing the Voronoi diagram in the plane is a consequence of the fact that the unbounded edges of the Voronoi diagram of a set $\mathcal{M}$ of points correspond to projections of the edges of the convex hull of $\mathcal{M}$. We also comment on this below. □

17.3 Delaunay complexes

17.3.1 Definition and connection with Voronoi diagrams

Given a set of n points $\mathcal{M} = \{M_1, \ldots, M_n\}$ in $\mathbb{E}^d$, we lift the points onto the paraboloid $\mathcal{P}$ to $\{\phi(M_1), \ldots, \phi(M_n)\}$, and consider the unbounded polytope $\mathcal{V}(\mathcal{M})$ that is the intersection of the n half-spaces that lie above the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$, where $\phi(M_i)^*$ is tangent to $\mathcal{P}$ at $\phi(M_i)$.

We denote by $\mathcal{D}(\mathcal{M})$ the convex hull of the points $\phi(M_1), \ldots, \phi(M_n)$ and a point O' on the x_{d+1}-axis, with $x_{d+1} > 0$ large enough so that the facial structure of

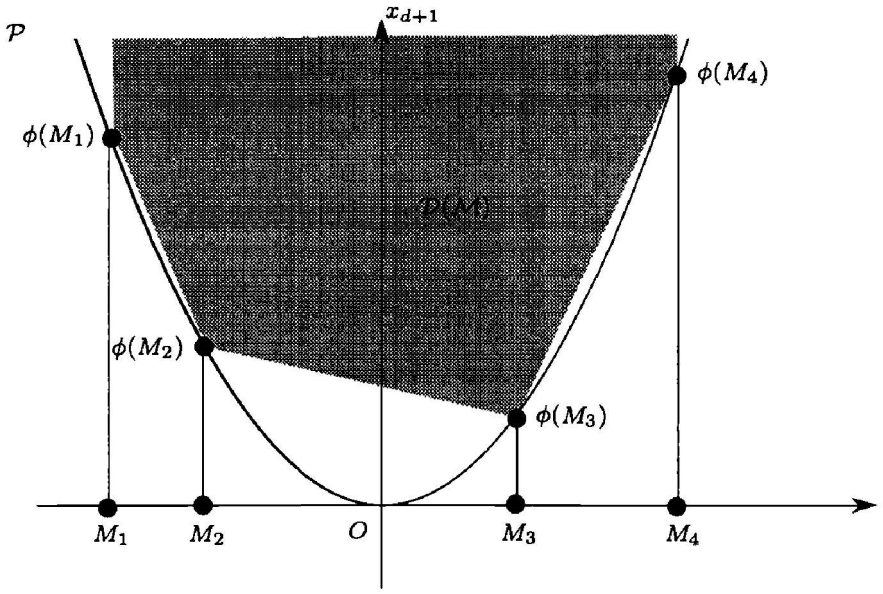

Figure 17.6. The polytope $\mathcal{D}(\mathcal{M})$.

that convex hull is stable as O' vanishes to infinity (see figure 17.6). The faces of $\mathcal{D}(\mathcal{M})$ that do not contain O' form the lower envelope of $conv(\phi(M_1), \ldots, \phi(M_n))$ (see also exercise 7.14). Their projections onto $\mathbb{E}^d$ form a complex whose vertices are exactly the M_i's. The domain of this complex is the projection of the convex hull of the $\phi(M_i)$'s: it is therefore the convex hull $conv(\mathcal{M})$ of the M_i's. This complex is called the Delaunay complex of $\mathcal{M}$ and is denoted by $\mathcal{D}el(\mathcal{M})$. For $k = 0, \ldots, d$, the k-faces of $\mathcal{D}el(\mathcal{M})$ are thus in one-to-one correspondence with the k-faces of $\mathcal{D}(\mathcal{M})$ that do not contain O'.

As shown in exercise 7.14, there exists a bijection between the faces of $\mathcal{V}(\mathcal{M})$ and the faces of $\mathcal{D}(\mathcal{M})$ that do not contain O'. This bijection maps the facet of $\mathcal{V}(\mathcal{M})$ containing $\phi(M_i)^*$ to the point $\phi(M_i)$. More generally, the k-faces of $\mathcal{V}(\mathcal{M})$ are in one-to-one correspondence with the $(d-k)$-faces of $\mathcal{D}(\mathcal{M})$ that do not contain O'. Moreover, this bijection reverses inclusion relationships.

Owing to theorem 17.2.5, the k-faces of $\mathcal{V}or(\mathcal{M})$ are also in bijection with the k-faces of the unbounded polytope $\mathcal{V}(\mathcal{M})$. So we have a bijection between the k-faces of $\mathcal{V}or(\mathcal{M})$ and the $(d-k)$-faces of $\mathcal{D}el(\mathcal{M})$ that reverses inclusion relationships. The Delaunay complex $\mathcal{D}el(\mathcal{M})$ is therefore dual to the Voronoi diagram $\mathcal{V}or(\mathcal{M})$.

Notice that the duality above maps a face of $\mathcal{V}or(\mathcal{M})$, formed by the points equidistant from m sites in $\mathcal{M}$, to the face of $\mathcal{D}el(\mathcal{M})$ that is the convex hull of these sites.

The preceding discussion leads to the following theorem:

Theorem 17.3.1 *The Delaunay complex of n points $M_1, \ldots, M_n$ in $\mathbb{E}^d$ is a complex dual to the Voronoi diagram. Its faces are obtained by projecting the faces of the lower envelope of the convex hull of the n points $\phi(M_1), \ldots, \phi(M_n)$, obtained by lifting the M_i's onto the paraboloid $\mathcal{P}$.*

The preceding theorem reduces the computation of the Delaunay complex of n points in $\mathbb{E}^d$ to the computation of the convex hull of n points in $\mathbb{E}^{d+1}$. All the convex hull algorithms described in this book, be they deterministic or randomized, static or dynamic, output-sensitive or not, therefore provide algorithms of the same kind that compute Delaunay complexes.

Theorem 17.3.1 also gives a lower bound on the complexity of the Delaunay complex. Indeed, exercise 7.11 exhibits polytopes in $\mathbb{E}^{d+1}$ with n vertices on the paraboloid, whose complexity is $\Theta(n^{\lceil d/2 \rceil})$. The same bound therefore applies to Delaunay complexes, and dually to Voronoi diagrams.

Corollary 17.3.2 *The Delaunay complex of n points in $\mathbb{E}^d$ can be computed in time $O(n \log n + n^{\lceil d/2 \rceil})$, and this is optimal in the worst case.*

17.3.2 Delaunay triangulations

Under L_2-general position assumptions, $\mathcal{V}(\mathcal{M})$ is a simple polytope, $\mathcal{D}(\mathcal{M})$ is a simplicial polytope, and $\mathcal{D}el(\mathcal{M})$ is a simplicial complex which we call in this case the *Delaunay triangulation* (see figure 17.7). If there is a subset $\mathcal{M}' \subset \mathcal{M}$ of $l > d+1$ co-spherical points and if the interior of the sphere circumscribed to $\mathcal{M}'$ (namely the sphere that passes through all the points in $\mathcal{M}'$) does not contain points in $\mathcal{M} \setminus \mathcal{M}'$, the Delaunay complex $\mathcal{D}(\mathcal{M})$ is not simplicial any more since $conv(\mathcal{M}')$ is a d-face of the Delaunay complex and it is not a simplex. Note however that this face may always be triangulated, and other non-simplicial faces of the complex may be triangulated as well. There are many ways to triangulate these faces, and any such triangulation is called a *Delaunay triangulation*. Henceforth, we denote by $\mathcal{D}et(\mathcal{M})$ any such triangulation.

17.3.3 Characteristic properties

The Delaunay complex has remarkable properties, all due to the fact that it is dual to the Voronoi diagram.

Theorem 17.3.3 *Let $\mathcal{M}$ be a set of n points $M_1, \ldots, M_n$ in $\mathbb{E}^d$. Any d-face in the Delaunay complex can be circumscribed by a sphere that passes through all its vertices, and whose interior contains no point in $\mathcal{M}$.*

Proof. Let us pick a d-face T of the Delaunay complex. Then T is the convex hull $T = conv(M_{i_0}, \ldots, M_{i_l})$ of l co-spherical points $M_{i_0}, \ldots, M_{i_l}$. (If the points are in

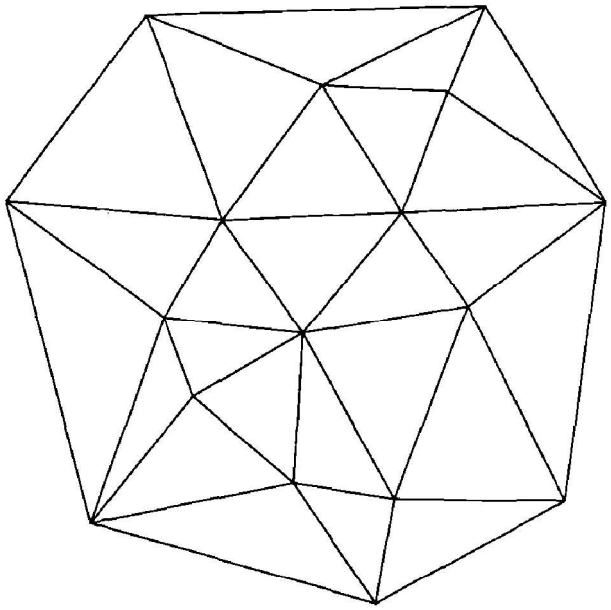

Figure 17.7. The Delaunay triangulation that corresponds to the Voronoi diagram shown in figure 17.1.

L_2-general position we have $l = d$.) The convex hull $conv(\phi(M_{i_0}), \ldots, \phi(M_{i_l}))$ is a d-face F of the convex hull $conv(\phi(M_1), \ldots, \phi(M_n))$, because of theorem 17.3.1. The intersection of the hyperplane $\underline{H}_F$ that supports F and of the paraboloid projects onto $\mathbb{E}^d$ as a sphere Σ circumscribed to $conv(M_{i_0}, \ldots, M_{i_l})$, and its center is the projection on $\mathbb{E}^d$ of the pole $\underline{H}_F^*$ of $\underline{H}_F$ (see lemma 17.2.2). $\underline{H}_F^*$ is a vertex of $\mathcal{V}(\mathcal{M})$, and more precisely is the intersection of the polar hyperplanes $\phi(M_{i_0}^*), \ldots, \phi(M_{i_l}^*)$. C is the vertex of the Voronoi diagram that is incident to the cells that correspond to the sites $M_{i_0}, \ldots, M_{i_l}$, and the interior of Σ cannot contain any other point in $\mathcal{M}$. □

Our next theorem extends this result into a necessary and sufficient condition for the convex hull of some points in $\mathcal{M}$ to be a face of the Delaunay complex of $\mathcal{M}$.

Theorem 17.3.4 *Let $\mathcal{M}$ be a set of points in $\mathbb{E}^d$, and $\mathcal{M}_k = \{M_{i_0}, \ldots, M_{i_k}\}$ be a subset of k points in $\mathcal{M}$. The convex hull of $\mathcal{M}_k$ is a face of the Delaunay complex if and only if there exists a $(d-1)$-sphere passing through $M_{i_0}, \ldots, M_{i_k}$ and such that no point in $\mathcal{M}$ belongs to its interior.*

Proof. The necessary condition immediately results from the preceding theorem and from the fact that a sphere circumscribed to a face is also circumscribed to its subfaces. Assume that there exists a $(d-1)$-sphere Σ that passes through $M_{i_0}, \ldots, M_{i_k}$ and whose interior contains no point in $\mathcal{M}$. Let $\underline{H}$ be the hyperplane $\phi(\Sigma)^*$ in $\mathbb{E}^{d+1}$. This hyperplane contains the points $\phi(M_{i_0}), \ldots, \phi(M_{i_k})$

and the half-space $\underline{H}^-$ lying below $\underline{H}$ does not contain points in $\phi(\mathcal{M})$ (according to lemma 17.2.4). Thus $\underline{H}$ is a hyperplane supporting $\mathcal{D}(\mathcal{M})$ along $conv(\phi(M_{i_0}), \ldots, \phi(M_{i_k}))$. Hence $conv(\phi(M_{i_0}) \ldots \phi(M_{i_k})) = \underline{H} \cap \mathcal{D}(\mathcal{M})$ is a face of $\mathcal{D}(\mathcal{M})$. It follows from theorem 17.3.1 that $M_{i_0}, \ldots, M_{i_k}$ is a face of the Delaunay complex of $\mathcal{M}$. □

Corollary 17.3.5 *Any Delaunay triangulation of a set $\mathcal{M}$ of points in $\mathbb{E}^d$ is such that the sphere circumscribed to any d-simplex in the triangulation contains no point of $\mathcal{M}$ in its interior. Conversely, any triangulation satisfying this property is a Delaunay triangulation.*

The next theorem now exhibits a local characterization of Delaunay triangulations that will be put to good use later on. Let us consider any Delaunay triangulation $\mathcal{T}(\mathcal{M})$ of a set $\mathcal{M}$ of points in $\mathbb{E}^d$ and let $S_1 = M_1 \ldots M_d M_{d+1}$ and $S_2 = M_1 \ldots M_d M_{d+2}$ be a pair of adjacent d-simplices in $\mathcal{T}(\mathcal{M})$ that share a common face $F = M_1 \ldots M_d$. The pair (S_1, S_2) is called *regular* if M_{d+1} does not belong to the interior of the sphere Σ_2 circumscribed to S_2. If the sphere Σ_1 circumscribed to S_1 differs from Σ_2, the regularity condition is equivalent to the property that M_{d+2} does not belong to the interior of Σ_1. Indeed, M_{d+1} does not belong to the interior of Σ_2 if and only if $\Sigma_2(M_{d+1}) > 0$. But the hyperplane H_F that supports F is the radical hyperplane of Σ_1 and Σ_2. Since $\Sigma_1(M_{d+1}) = 0$, the half-space bounded by H_F that contains M_{d+1} (resp. M_{d+2}) consists of the points whose power with respect to Σ_1 is smaller (resp. greater) than their power with respect to Σ_2, and therefore

$$\Sigma_1(M_{d+2}) > \Sigma_2(M_{d+2}) = 0,$$

which proves that M_{d+2} does not belong to the interior of Σ_1.

Theorem 17.3.6 *Consider a triangulation $\mathcal{T}(\mathcal{M})$ of a set $\mathcal{M}$ of points in $\mathbb{E}^d$. Then $\mathcal{T}(\mathcal{M})$ is a Delaunay triangulation if and only if all the pairs of adjacent d-simplices in $\mathcal{T}(\mathcal{M})$ are regular.*

Proof. That the condition is necessary is a consequence of theorem 17.3.3. We must now show that it suffices. To alleviate the notation, we denote by $\phi(S)$ the k-simplex in $\mathbb{E}^{d+1}$ whose vertices are the images of the vertices of a k-simplex S in $\mathbb{E}^d$, and by $\mathcal{C}$ the union of the $\phi(S)$'s for all the faces S of the Delaunay triangulation $\mathcal{T}(\mathcal{M})$. The proof consists of proving that $\mathcal{C}$ is the graph of a convex real-valued function over the convex hull $conv(\mathcal{M})$.

As above, we consider two adjacent d-simplices $S_1 = M_1 \ldots M_d M_{d+1}$ and $S_2 = M_1 \ldots M_d M_{d+2}$ in $\mathcal{T}(\mathcal{M})$ that share a common face $F = M_1 \ldots M_d$. We denote

by Σ_1 and Σ_2 the spheres circumscribed to S_1 and S_2. Owing to lemma 17.2.4, the regularity condition is equivalent to $\phi(M_{d+1}) \in \phi(\Sigma_2)^{*+}$ and also to $\phi(M_{d+2}) \in \phi(\Sigma_1)^{*+}$ because of the discussion above. Therefore, if the pair (S_1, S_2) is regular, then the $(d-1)$-face $\phi(F)$ is locally convex, meaning that there is a hyperplane that contains $\phi(F)$ such that $\phi(S_1)$ and $\phi(S_2)$ belong to the half-space lying above this hyperplane. This is true for any $(d-1)$-face of $\mathcal{C}$ incident to two d-faces, and so $\mathcal{C}$ is locally convex at any point. Moreover, $\mathcal{C}$ is defined over a convex subset of $\mathbb{E}^d$, namely the convex hull of $\mathcal{M}$. Therefore, $\mathcal{C}$ is convex and is the lower envelope of the polytope $\mathcal{D}(\mathcal{M})$, which proves that $\mathcal{T}(\mathcal{M})$ is a Delaunay triangulation of $\mathcal{M}$. □

17.3.4 Optimality of Delaunay triangulations

As we have seen in chapter 11, there exist several ways to triangulate a set of points. Some are not very interesting in practice, and in many applications certain criteria must be optimized, and an optimal triangulation is desirable. There are several ways to define optimality. In this section, we show that Delaunay triangulations maximize two criteria, compactness and equiangularity.

Compactness

The preceding theorem was concerned with spheres circumscribed to simplices in the triangulation. The next theorem considers the *smallest enclosing sphere* for each simplex S: this sphere is the circumscribed sphere of S if the center of the latter belongs to S, or otherwise is a sphere centered on some k-face ($k < d$) of S and passes through the $k+1$ centers of this face.

As before, we consider a set $\mathcal{M}$ of points in $\mathbb{E}^d$ and $\mathcal{T}(\mathcal{M})$ a triangulation of $\mathcal{M}$. To $\mathcal{T}(\mathcal{M})$ corresponds a function $\Sigma_{\mathcal{T}}(X)$ defined over $conv(\mathcal{M})$ as the power of a point X with respect to the sphere Σ circumscribing any d-simplex of $\mathcal{T}(\mathcal{M})$ that contains X. By the results of subsection 17.2.6, $\Sigma_{\mathcal{T}}(X)$ is well-defined when X belongs to several cells.

Lemma 17.3.7 *Let $\mathcal{D}et(\mathcal{M})$ be a Delaunay triangulation of $\mathcal{M}$ and $\mathcal{T}(\mathcal{M})$ be any other triangulation of $\mathcal{M}$. Then*

$$\forall X \in conv(\mathcal{M}),\ \Sigma_{\mathcal{D}et}(X) \geq \Sigma_{\mathcal{T}}(X).$$

Proof. Consider a d-simplex T in $\mathcal{T}(\mathcal{M})$ that contains X, Σ its circumscribed sphere, and $\phi(T)$ the d-simplex of $\mathbb{E}^{d+1}$ whose vertices are the images under ϕ of the vertices of T. (Recall that these vertices are obtained by lifting the vertices of T onto the paraboloid $\mathcal{P}$.) Lemma 17.2.3 shows that $\Sigma_{\mathcal{T}}(X)$ is the signed vertical distance (here negative) from $\phi(\Sigma)^*$ to $\phi(X)$. Notice that $\phi(\Sigma)^*$ is the affine hull

of $\phi(T)$. For a given X, this signed vertical distance is maximized when $\phi(T)$ is a face of the convex hull of $\phi(\mathcal{M})$ in $\mathbb{E}^{d+1}$: in other words, when T is a simplex of a Delaunay triangulation of $\mathcal{M}$. □

Lemma 17.3.8 *If T is a d-simplex and if Σ_T is its circumscribed sphere, then*

$$\min_{X \in T} \Sigma_T(X) = \Sigma_T(C'_T) = -{r'_T}^2$$

where C'_T and r'_T are respectively the center and the radius of the smallest sphere enclosing T.

Proof. Let Σ_T be the sphere circumscribed to T, C_T its center, and r_T its radius. Then

$$\Sigma_T(X) = XC_T^2 - r_T^2$$

is minimized when $X = C_T$ and is therefore greater than $-r_T^2$. If C_T is contained in T, the smallest enclosing sphere of T is Σ_T, hence $r'_T = r_T$ and the lemma is trivial. Otherwise, the smallest enclosing sphere of T is centered on a k-face ($k < d$), namely the face F such that the orthogonal projection of C_T onto the plane that supports F falls inside F. The radius r'_T of this sphere is that of the $(k-1)$-sphere circumscribed to F. Its center C'_T minimizes the value of $XC_T \cdot XC_T$ when $X \in T$. Pythagoras' theorem then shows that

$$C_T{C'_T}^2 + {r'_T}^2 = r_T^2,$$

which finishes the proof. □

Let $\mathcal{T}(\mathcal{M})$ be any triangulation of a set $\mathcal{M}$ of points in $\mathbb{E}^d$. For each simplex T in $\mathcal{T}(\mathcal{M})$, we let r'_T denote the smallest radius of a sphere that encloses T, and the *maximum min-containment radius* of $\mathcal{T}(\mathcal{M})$ is defined by

$$C(\mathcal{T}(\mathcal{M})) = \max_{T \in \mathcal{T}(\mathcal{M})} r'_T.$$

The *most compact* triangulations are then defined as the triangulations that minimize the maximum min-containment radius.

Theorem 17.3.9 *Delaunay triangulations are the most compact among all the triangulations of $\mathcal{M}$.*

Note that since the maximum min-containment radius $C(\mathcal{T}(\mathcal{M}))$ is defined only by the simplices T of $\mathcal{T}(\mathcal{M})$ such that $C(\mathcal{T}(\mathcal{M})) = r'_T$, triangulations other than Delaunay triangulations might also be most compact among the triangulations of $\mathcal{M}$.

Proof. Let $\mathcal{T}(\mathcal{M})$ be any triangulation of $\mathcal{M}$ and let $\mathcal{D}et(\mathcal{M})$ be any Delaunay

triangulation of $\mathcal{M}$. We denote by $X_{\mathcal{T}}$ the point X that minimizes $\Sigma_{\mathcal{T}}(X)$ and by $X_{\mathcal{D}et}$ the point X that minimizes $\Sigma_{\mathcal{D}et}(X)$. From lemma 17.3.8, we know that $X_{\mathcal{T}}$ is the center of the smallest sphere that encloses the simplex in $\mathcal{T}(\mathcal{M})$ which contains $X_{\mathcal{T}}$. We denote its radius by $r'_{\mathcal{T}}$. Likewise, $X_{\mathcal{D}et}$ is the center of the smallest sphere that encloses the simplex in $\mathcal{D}et(\mathcal{M})$ that contains $X_{\mathcal{T}}$ and its radius is denoted by $r'_{\mathcal{D}et}$. The maximum min-containment radius of $\mathcal{T}(\mathcal{M})$ equals $r'_{\mathcal{T}}$ and that of $\mathcal{D}et$ equals $r'_{\mathcal{D}et}$. Using lemmas 17.3.7 and 17.3.8, we obtain

$$\Sigma_{\mathcal{T}}(X_{\mathcal{T}}) = -{r'_{\mathcal{T}}}^2 \leq \Sigma_{\mathcal{T}}(X_{\mathcal{D}et}) \leq \Sigma_{\mathcal{D}et}(X_{\mathcal{D}et}) = -{r'_{\mathcal{D}et}}^2.$$

□

Equiangularity ($d = 2$)

We now restrict the discussion to triangulations of a set of points in the plane. Given a triangulation $\mathcal{T}(\mathcal{M})$ of a set $\mathcal{M}$ of n points in the plane, we define its *angle vector* as the vector $Q(\mathcal{T}(\mathcal{M})) = (\alpha_1, \ldots, \alpha_{3t})$ where the α_i's are the angles of the t triangles of $\mathcal{T}(\mathcal{M})$ sorted by increasing value. We know that $\sum_{i=1}^{3t} \alpha_i = t\pi$. Note that a triangulation that maximizes the angle vector for the lexicographic order also maximizes the smallest of its angles. Such a triangulation is called *globally equiangular.*

Theorem 17.3.10 *A globally equiangular triangulation of a set $\mathcal{M}$ of points in the plane is always a Delaunay triangulation.*

Proof. We must prove that, among all the triangulations of $\mathcal{M}$, the ones that maximize the angle vector for the lexicographic order are always Delaunay triangulations. Let us thus consider two triangles $T_1 = ABC$ and $T_2 = BCD$ in some triangulation $\mathcal{T}(\mathcal{M})$, such that the union of T_1 and T_2 is a strictly convex quadrilateral $\mathcal{Q}$. (This means that A, B, C, and D are all vertices of the convex hull $conv(A, B, C, D)$.) In order to increase the equiangularity, we can flip the diagonal as follows (shown in figure 17.8). If the triangles $T'_1 = ABD$ and $T'_2 = ACD$ are such that $Q(T'_1, T'_2) > Q(T_1, T_2)$, then replace $\mathcal{T}(\mathcal{M})$ by a triangulation $\mathcal{T}^1(\mathcal{M})$ which contains T'_1 and T'_2 instead of T_1 and T_2.

The previous rule may be dubbed a *regularization* rule since it transforms a pair of adjacent triangles into a regular pair of triangles: if the two triangles do not form a convex quadrilateral, then the pair is obviously regular, and the rule does not apply; otherwise, $T_1 \cup T_2$ is convex and the pair is transformed into a regular pair. Indeed, let Σ_1 and Σ_2 be the circles circumscribed to T_1 and T_2. We will show that the diagonal AD is flipped if and only if D is contained inside the circle Σ_1. Let α, β, γ, and δ be the angles at the vertices of the quadrilateral $ABCD$, α, β_1, and γ_1 the angles at the vertices of T_1, and β_2, γ_2, and δ the angles

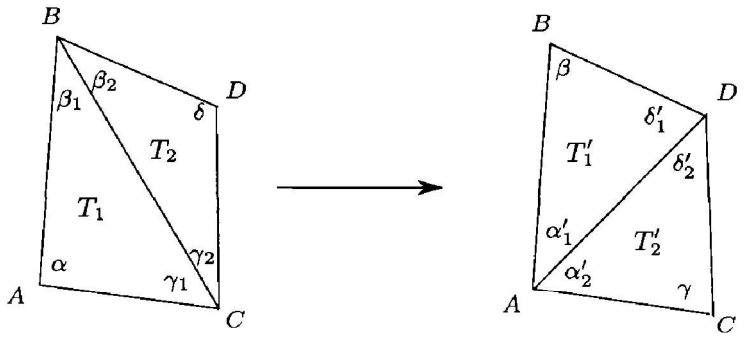

Figure 17.8. Flipping a diagonal to increase the equiangularity.

at the vertices of T_2. Moreover, we denote by α_1', β, and δ_1' the angles at the vertices of T_1', and α_2', γ, and δ_2' the angles at the vertices of T_2'. The situation is depicted in figure 17.8. If α is the smallest angle in T_1 and T_2, then the diagonal is not flipped. But then

$$\delta = \pi - \beta_2 - \gamma_2 < \pi - \alpha,$$

so that $\alpha + \delta < \pi$, which shows that A is not contained inside Σ_2, and this implies that D is not contained inside Σ_1. The situation is entirely symmetric when the smallest angle is δ. When the smallest angle is β_1, then we flip the diagonal only if δ_2' is greater than β_1, which only happens when D is contained inside Σ_1. Of course, the cases when the smallest angle is γ_1, β_2, or γ_2 are entirely similar, so we have shown that the diagonal is flipped if and only if it transforms the irregular pair (T_1, T_2) into a regular pair (T_1', T_2').

Clearly, after a flip we have $Q(\mathcal{T}^1(\mathcal{M})) > Q(\mathcal{T}(\mathcal{M}))$. Flipping the edges whenever possible progressively increases the angle vector of the triangulation. Since there are only a finite number of triangulations, this process eventually reaches a triangulation that has only regular pairs of adjacent triangles. This triangulation is a Delaunay triangulation as is shown by theorem 17.3.6. □

Note that this local regularization always leads to a Delaunay triangulation. When the points are in L_2-general position, there is only one Delaunay triangulation: the result of the procedure described above therefore does not depend on the starting configuration, nor on the order chosen to flip the diagonals.

When the points are not in L_2-general position, however, the theorem above shows that flipping diagonals only reaches a Delaunay triangulation. Yet there are several Delaunay triangulations, which may not all have the same angle vectors. Still, there is an algorithm that can reach a globally equiangular triangulation (see the bibliographical notes).

17.4 Higher-order Voronoi diagrams

In this section, we define Voronoi diagrams of order k and show the connection between these diagrams and the faces at level k in a hyperplane arrangement in $\mathbb{E}^{d+1}$. As usual, the Euclidean space $\mathbb{E}^d$ of dimension d is embedded in $\mathbb{E}^{d+1}$ as the hyperplane $x_{d+1} = 0$, and $\phi(M)^*$ denotes the hyperplane in $\mathbb{E}^{d+1}$ that is tangent to the paraboloid $\mathcal{P}$ at the point $\phi(M)$ obtained by lifting M vertically onto the paraboloid.

In section 17.2, we established the connection between the Voronoi diagram of a set $\mathcal{M}$ of n points $M_1, \ldots, M_n$ in $\mathbb{E}^d$ and the polytope in $\mathbb{E}^{d+1}$ that is the intersection of the n half-spaces $\phi(M_i)^{*+}$ that lie above the hyperplanes $\phi(M_i)$. Equivalently, $\mathcal{V}(\mathcal{M})$ is the cell at level 0 in the arrangement $\mathcal{A}$ of the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$, if the reference point is on the x_{d+1}-axis, sufficiently high so that it is above all the hyperplanes. Let us recall that a point is at level k in $\mathcal{A}$ if it belongs to exactly k open half-spaces $\phi(M_{i_1})^{*-}, \ldots, \phi(M_{i_k})^{*-}$, such that each $\phi(M_{i_j})^{*-}$ is bounded by $\phi(M_{i_j})^*$ and does not contain the reference point (see section 14.5).

It is tempting to consider the cells at levels $k > 0$. We define below a cell complex that spans $\mathbb{E}^d$, called the *Voronoi diagram of order k* of $\mathcal{M}$, and show in theorem 17.4.1 that the cells of this complex are the non-overlapping projections onto $\mathbb{E}^d$ of the cells at level k in the arrangement $\mathcal{A}$.

Let $\mathcal{M}_k$ be a subset of size k of $\mathcal{M}$. The Voronoi region of $\mathcal{M}_k$ is the polytope $V_k(\mathcal{M}_k)$ of the points in $\mathbb{E}^d$ that are closer to all the sites in $\mathcal{M}_k$ than to any other site in $\mathcal{M} \setminus \mathcal{M}_k$. Formally,

$$V_k(\mathcal{M}_k) = \{X : \forall M_i \in \mathcal{M}_k,\ \forall M_j \in \mathcal{M} \setminus \mathcal{M}_k,\ \|XM_i\| \leq \|XM_j\|\}.$$

Let us consider all the subsets of size k of $\mathcal{M}$ whose Voronoi regions are not empty. As proved in the theorem below, these polytopes and their faces form a d-complex whose domain is $\mathbb{E}^d$. This complex is called the *Voronoi diagram of order k* of $\mathcal{M}$ (see figures 17.9 and 17.10). It is denoted by $\mathcal{V}or_k(\mathcal{M})$. When $k = 1$, we recognize the definition of the usual Voronoi diagram.

Theorem 17.4.1 *The Voronoi diagram $\mathcal{V}or_k(\mathcal{M})$ of order k of a set $\mathcal{M} = \{M_1, \ldots, M_n\}$ of n points in $\mathbb{E}^d$ is a cell complex of dimension d in $\mathbb{E}^d$. The cells of this complex correspond to the cells at level k in the arrangement $\mathcal{A}$ of the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$ in $\mathbb{E}^{d+1}$, when the reference point is on the x_{d+1}-axis above all the hyperplanes. A cell of $\mathcal{V}or_k(\mathcal{M})$ is obtained by projecting vertically onto $\mathbb{E}^d$ the corresponding cell in $\mathcal{A}$. The l-faces of $\mathcal{V}or_k(\mathcal{M})$ are obtained by projecting the l-faces common to cells at level k in $\mathcal{A}$.*

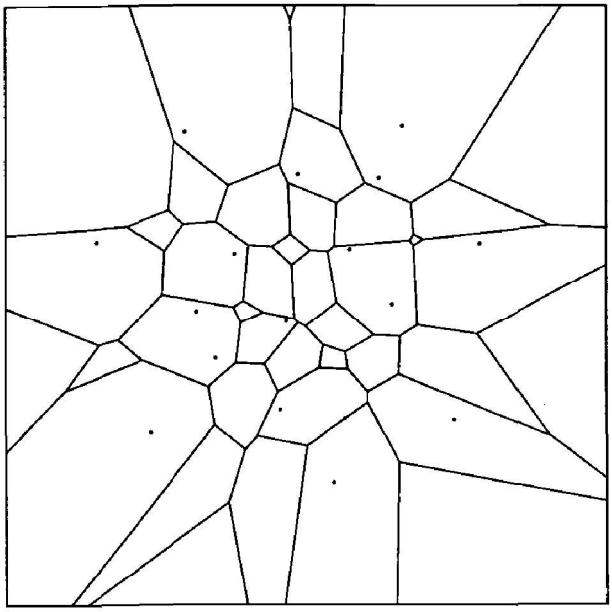

Figure 17.9. The Voronoi diagram of order 2 of the points in figure 17.1.

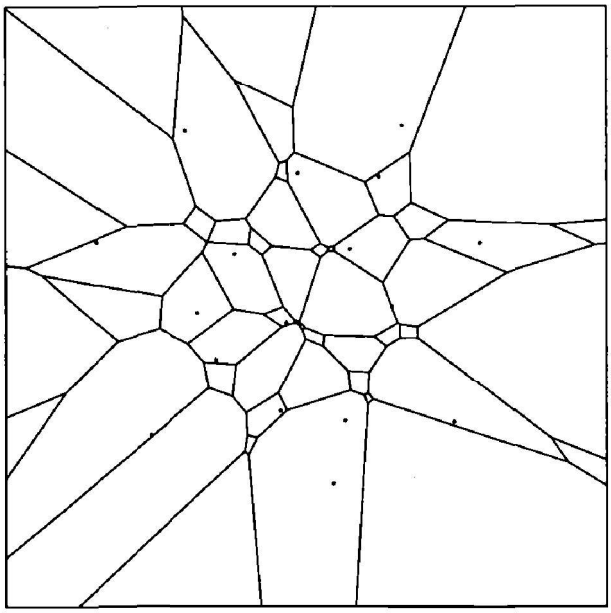

Figure 17.10. The Voronoi diagram of order 3 of the points in figure 17.1.

Proof. The proof relies on lemma 17.2.4. A sphere in $\mathbb{E}^d$ whose interior contains k points is mapped by ϕ to a point at level k in the arrangement $\mathcal{A}$ of the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$.

More precisely, X belongs to the cell $V_k(\mathcal{M}_k)$ in the Voronoi diagram of order k, if and only if X is the center of a sphere Σ whose interior contains the points

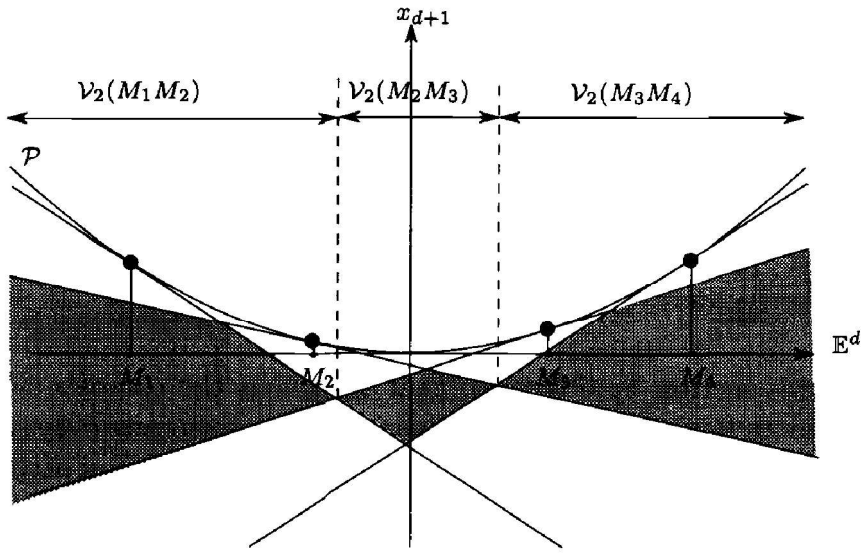

Figure 17.11. The Voronoi diagram of order 2 is obtained by projecting the cells at level 2 in $\mathcal{V}(\mathcal{M})$.

of $\mathcal{M}_k$ and only those. Then $\phi(\Sigma)$ belongs to the k closed half-spaces below the hyperplanes $\phi(M_j)^*$ for the M_j's in $\mathcal{M}_k$, and only to those half-spaces. The cells of the Voronoi diagram of order k are obtained by projecting vertically the cells at level k of $\mathcal{A}$ (see figure 17.11).

It is easily verified that any vertical line intersects at least one cell at level k in $\mathcal{A}$ and does not intersect the interior of more than one cell at level k. It follows that the l-faces, for $l < d$, of the Voronoi diagram of order k are obtained by vertically projecting the l-faces common to several cells at level k. If the M_i's are in L_2-general position, then the hyperplanes $\phi(M_1)^*, \ldots, \phi(M_n)^*$ are in general position. In that case, it was shown in section 14.5 that the cells of $\mathcal{A}$ that contain an l-face F at level k have levels that vary between k and $k+d+1-l$. Among those, there is only one cell at level k and one cell at level $k+d+1-l$, and several cells at levels $k < j < k+d+1-l$. It follows that the vertical projection of F is an l-face of the Voronoi diagrams of orders $k+1, k+2, \ldots, k+d-l$. □

Having computed the Voronoi diagram of order k of the sites, looking for the k nearest sites of any point X in $\mathbb{E}^d$ can be performed by finding the cell of the diagram that contains X (see exercises 17.2 and 17.4).

It follows from the construction that the total complexity of the Voronoi diagrams of all orders k, $1 \leq k \leq n-1$, is $O(n^{d+1})$: indeed it is exactly the complexity of the arrangement of the hyperplanes in $\mathbb{E}^{d+1}$. Moreover, these diagrams can be computed in time $O(n^{d+1})$ (see theorem 14.4.4). An upper bound on the complexity of the Voronoi diagrams $\mathcal{V}or_1(\mathcal{M}), \ldots, \mathcal{V}or_k(\mathcal{M})$ of orders between 1 and k is provided by theorem 14.5.1, which bounds the complexity

of the first k levels in the arrangement of n hyperplanes in $\mathbb{E}^{d+1}$. Moreover, $\mathcal{V}or_{\leq k}(\mathcal{M}) = \{\mathcal{V}or_1(\mathcal{M}), \ldots, \mathcal{V}or_k(\mathcal{M})\}$ can be computed by computing the first k levels in this arrangement (see theorem 14.5.3 and exercise 14.22) and projecting back onto $\mathbb{E}^d$.

Theorem 17.4.2 *The overall complexity of the first k Voronoi diagrams of a set of n points in $\mathbb{E}^d$ is $O(n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil})$. These k diagrams may be computed in time $O(n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil})$ if $d \geq 3$, or $O(nk^2 \log \frac{n}{k})$ if $d = 2$.*

Let us close this section by observing that $\mathcal{V}or_{n-1}$ is the complex whose cells consist of the points further from a particular site than from any other site. This is why this diagram is sometimes called the *furthest-point Voronoi diagram*. The vertices of this diagram are the centers of spheres circumscribed to $d+1$ sites and whose interiors contain all the other sites. Its cells are all unbounded. The furthest-point diagram can be obtained by computing the intersection of the n *lower* half-spaces bounded by the hyperplanes $\phi(M_i)^*$, $i = 1, \ldots, n$.

17.5 Exercises

Exercise 17.1 (Shortest edge) Denote by $\mathcal{F}$ and $\mathcal{G}$ two finite sets of points in $\mathbb{E}^d$. Show that the shortest edge that connects a point in $\mathcal{F}$ to a point in $\mathcal{G}$ is an edge of the Delaunay triangulation of $\mathcal{F} \cup \mathcal{G}$. From this, conclude that each point is adjacent to its nearest neighbor in the Delaunay triangulation.

Exercise 17.2 (Nearest neighbor in the plane) Show that, given the Voronoi diagram of a set $\mathcal{M}$ of points in the plane, it may be preprocessed in linear time to answer nearest neighbor queries (that is, find the nearest site to a query point) in logarithmic time. Same question for the set of k nearest neighbors (k fixed).

Hint: One may use the result of exercise 12.2.

Exercise 17.3 (On-line nearest neighbor) We place point sites in the plane and we want to maintain a data-structure on-line so as to answer nearest neighbor queries on the current set of sites (that is, find the nearest site to a query point). Devise a structure that stores n sites using storage $O(n)$, such that under the assumption that the points are inserted in a random order, the expected time needed to insert a new site is $O(\log n)$, and that answers any query in expected time $O(\log^2 n)$.

Hint: One may build a two-level data-structure in the following way. The first level corresponds to a triangulation of the Voronoi diagram, obtained by connecting a site to all the vertices of its Voronoi region. Build an influence graph for the regions defined as the triangles in this triangulation (it is a variant of the influence graph used in exercise 17.10). Each triangle points to the site that kills it, and all the triangles created after the insertion of a site are sorted in polar angle around this site and stored into an array: this

is the second level of the data structure. Show that the query point belongs to $O(\log n)$ triangles on the average, and hence that only $O(\log n)$ binary searches are performed in the arrays of the second level.

Exercise 17.4 (Nearest neighbor) Consider a set of n point sites in $\mathbb{E}^d$. Explain how to design a data structure of size $O(n^{\lceil d/2 \rceil + \varepsilon})$ that allows the nearest site to any query point P to be found in logarithmic time.

Hint: Use the solution of exercise 14.12 in $\mathbb{E}^{d+1}$. The exponent ε can be removed, at the cost of increasing the query time to $O(\log^2 n)$ (see the bibliographical notes at the end of chapter 14).

Exercise 17.5 (Union of balls) Use lemma 17.2.4 to reduce the problem of computing the union of n balls in $\mathbb{E}^d$ to that of computing the intersection of the paraboloid $\mathcal{P}$ with a polytope in $\mathbb{E}^{d+1}$. Conclude that the complexity of the union of n balls is $O(n^{\lceil \frac{d}{2} \rceil})$. Devise an algorithm that computes the union of n balls in expected time $\Theta(n \log n + n^{\lceil \frac{d}{2} \rceil})$.

Exercise 17.6 (Intersection of balls) The results of exercise 17.5 are also valid for the intersection of n balls in $\mathbb{E}^d$. In $\mathbb{E}^3$, show that if the balls have same radius, the complexity of the intersection is only $O(n)$ and propose an algorithm that computes this intersection in expected time $\Theta(n \log n)$.

Hint: Show that each face of the intersection is "convex", meaning that given any two points in any face, there is an arc of a great circle joining these points which is entirely contained in that face; then use Euler's relation. For the algorithm, use a variant of the randomized incremental algorithm of section 8.3.

Exercise 17.7 (Minimum enclosing ball) Show that the center of the smallest ball whose interior contains a set $\mathcal{M}$ of points in $\mathbb{E}^2$ is either a vertex of the furthest-point Voronoi diagram (of order $n-1$) of $\mathcal{M}$, or else the intersection of an edge of this diagram (on the perpendicular bisector of two sites A and B) with the edge AB.

Exercise 17.8 (Centered triangulation) Consider a triangulation $\mathcal{T}(\mathcal{M})$ of a set $\mathcal{M}$ of points in $\mathbb{E}^d$. Show that, if each simplex in $\mathcal{T}(\mathcal{M})$ contains the center of its circumscribed sphere, then $\mathcal{T}(\mathcal{M})$ is a Delaunay triangulation of $\mathcal{M}$. Construct a counter-example for the converse.

Hint: Show that any adjacent pair of d-simplices is regular.

Exercise 17.9 (Non-optimality of the Delaunay triangulation) Construct a set of points in the plane whose Delaunay triangulation does not minimize the greatest angle among all the possible triangulations. Same question to show that the Delaunay triangulation does not necessarily minimize the length of the longest edge, nor the total length of the triangulation (sum of the lengths of the edges).

Exercise 17.10 (Incremental algorithm) Let $\mathcal{D}el(\mathcal{M})$ be the Delaunay triangulation of a set $\mathcal{M}$ of points in L_2-general position in $\mathbb{E}^d$. Let A be a point of $\mathbb{E}^d$ distinct from the points in $\mathcal{M}$, let $\mathcal{S}$ be the sub-complex formed by the d simplices in $\mathcal{D}el(\mathcal{M})$ whose circumscribed spheres contain A in their interior, and let $\mathcal{F}$ be the set of $(d-1)$-faces on the boundary of $\mathcal{S}$. Show that if A belongs to the convex hull of $\mathcal{M}$, then the d-simplices in $\mathcal{D}el(\mathcal{M} \cup \{A\})$ are exactly the simplices in $\mathcal{D}el(\mathcal{M})$ that do not belong to $\mathcal{S}$ and the simplices $conv(A, F)$, $F \in \mathcal{F}$. Generalize this result to the case when A is outside the convex hull of $\mathcal{M}$, and derive an incremental algorithm that computes the Delaunay triangulation. Show that this algorithm runs in time $\Theta(n^{\lceil \frac{d}{2} \rceil + 1})$ in the worst case. Show that if the points are inserted in random order, then the algorithm runs in expected time $O(n \log n + n^{\lceil \frac{d}{2} \rceil})$, which is optimal. Devise a dynamic algorithm that also allows points to be removed.

Hint: Use a randomized algorithm with an influence graph. Objects are sites, regions are the balls circumscribed to $d+1$ sites, and an object conflicts with a region if it belongs to that region. Show that the ball circumscribed to any new simplex $S = conv(A, F)$ is contained in the union of the two balls circumscribed to T and V, the two d-simplices that share the common facet F. Build an influence graph in which each node is the child of only two nodes, namely the node corresponding to F is the child of the nodes corresponding to T and V. The number of children of a node is not bounded, but the analysis can be carried out using biregions (see exercise 5.7).

Exercise 17.11 (Flipping the diagonals) Devise an incremental algorithm to compute the Delaunay triangulation of points in the plane which, at each step, connects the new point to the edges of the triangle that contains it, and then regularizes the triangulation as in the proof of theorem 17.3.10. Show that, if the points are inserted in random order, the algorithm can be made to run in expected time $O(n \log n)$, which is optimal.

Exercise 17.12 (Flipping in higher dimensions) Generalize the local regularization rule introduced in the proof of theorem 17.3.10 to the triangulation of point sites in $\mathbb{E}^3$ and in higher dimensional spaces. Show that this does not always result in the Delaunay triangulation of the points, in contrast with the planar case.

Hint: As was done for planar triangulations (proof of theorem 17.3.10), local regularization in $\mathbb{E}^3$ corresponds to replacing the upper facets of a simplex in $\mathbb{E}^4$ by its lower facets. A simplex in $\mathbb{E}^4$ having five facets, local regularization in $\mathbb{E}^3$ leads to replacing two adjacent tetrahedra T_1 and T_2 by three tetrahedra T_3, T_4, and T_5 that are pairwise adjacent (and have the same vertices as T_1 and T_2), or the converse. Show that the local regularization rule cannot always be applied even though the triangulation is not regular everywhere.

Exercise 17.13 (Flipping in higher dimensions) Show that if one adds a new point P to a Delaunay triangulation $\mathcal{D}et(\mathcal{M})$ of a set $\mathcal{M}$ of points in $\mathbb{E}^d$, the Delaunay triangulation $\mathcal{D}et(\mathcal{M} \cup \{P\})$ can be obtained by splitting the simplex of $\mathcal{D}et(\mathcal{M})$ that contains P into $d+1$ new simplices, and then applying the generalized local regularization of exercise 17.12. Show that if the n points in $\mathcal{M}$ are inserted in a random order, this incremental algorithm computes $\mathcal{D}et(\mathcal{M})$ in expected time $O(n \log n + n^{\lceil \frac{d}{2} \rceil})$, which is optimal.

Hint: As in exercise 17.12, each regularization in $\mathbb{E}^d$ corresponds to replacing the upper facets of a $(d+1)$-simplex in $\mathbb{E}^{d+1}$ by its lower facets. Show that any $(d+1)$-simplex S involved in any step of the local regularization has P as a vertex and that the convex hull of its vertices other than P is a d-simplex of $\mathcal{D}et(\mathcal{M})$ that is destroyed by the local regularization.

Exercise 17.14 (Complexity of the Voronoi diagram of order k) Show that the complexity of the Voronoi diagram of order k of n points in $\mathbb{E}^2$ is always $O(k(n-k))$ and can be $\Omega(k(n-k))$ in the worst case.

Exercise 17.15 (Higher-order Voronoi diagrams and polytopes) Let $\mathcal{M}$ be a set of n points $M_1, \ldots, M_n$ in $\mathbb{E}^d$. With each subset $\mathcal{M}_k = \{M_{i_1}, \ldots, M_{i_k}\}$ of size k of $\mathcal{M}$, we associate its center of gravity $G(\mathcal{M}_k) = \frac{1}{k}\sum_{j=1}^{k} M_{i_j}$, and the real number $\sigma(\mathcal{M}_k) = \frac{1}{k}\sum_{j=1}^{k} M_{i_j}^2$. Show that the Voronoi diagram of order k of $\mathcal{M}$ is the projection of the polytope in $\mathbb{E}^d$ defined as the intersection of the half-spaces lying above the hyperplanes polar to the points $(G(\mathcal{M}_k), \sigma(\mathcal{M}_k))$, for all subsets $\mathcal{M}_k$ of size k of $\mathcal{M}$.

Hint: From the fact that

$$\frac{1}{k}\sum_{j=1}^{k}(X - M_{i_j}) \cdot (X - M_{i_j}) = X^2 - \frac{2}{k}\sum_{j=1}^{k} M_{i_j} \cdot X + \frac{1}{k}\sum_{j=1}^{k} M_{i_j}^2,$$

we infer that the k nearest neighbors of X are the points in $\mathcal{M}_k$ if and only $\mathcal{M}_k$ is the subset for which X has the smallest power with respect to the sphere centered at $G(\mathcal{M}_k)$ and whose power with respect to the origin is $\sigma(\mathcal{M}_k)$.

Exercise 17.16 (Euclidean minimum spanning tree) Consider a set $\mathcal{M}$ of n points in L_2-general position in $\mathbb{E}^d$. A *Euclidean minimum spanning tree*, or EMST for short, is a tree whose nodes are the points in $\mathcal{M}$ and whose total edge length is minimal. Show that such a tree is a subgraph of the Delaunay triangulation of $\mathcal{M}$. For the planar case, show that an EMST can be computed in time $O(n \log n)$. Consider the case where the set of points is not in L_2-general position any more.

Hint: Show that the following greedy algorithm produces a minimum spanning tree. Denote by $\mathcal{A}$ the set of points of $\mathcal{M}$ that are already connected to the current tree. The greedy algorithm picks the shortest segment that does not induce a cycle in the current subtree. This edge connects a point of $\mathcal{A}$ to a point of $\mathcal{M} \setminus \mathcal{A}$. The latter point is added to $\mathcal{A}$, the edge is added to the tree, and so on until the tree spans $\mathcal{M}$. Show that this yields an EMST, even if the points are not in L_2-general position. Exercise 17.1 shows that it can be completed into a Delaunay triangulation of $\mathcal{M}$. Explain how to make the algorithm run in time $O(n \log n)$.

17.6 Bibliographical notes

Voronoi diagrams have been used for a long time and in various disguises. Voronoi, a Russian mathematician of the early twentieth century, was the first to give them a precise definition and study them for their own sake, but they had already been used by

Wigner and Steiz in crystallography, by Thiessen in geography, and even by Descartes in astronomy. Delaunay [78] established most of the fundamental properties about the triangulation that bears his name.

More recently, the connection between Voronoi diagrams and polytopes was discovered by Brown [38] and by Edelsbrunner and Seidel [96]. The presentation in this book is based on the works of Boissonnat, Cérézo, Devillers, and Teillaud [26] and of Devillers, Meiser, and Teillaud [82]. The connection with polytopes answers exercise 17.4.

The optimality of the Delaunay triangulation was established by Rajan [193] for the compactness, and by Lawson [144] for the equiangularity. Mount and Saalfeld [170] have proposed an algorithm to compute a globally equiangular triangulation when the points are not in L_2-general position. In the context of approximating surfaces by piecewise linear patches, controlling the equiangularity serves to control the quality of the approximation even though it is more profitable to minimize the greatest angle as was shown by Nielson and Franke [179]. The Delaunay triangulation does not generally minimize the greatest angle, nor the total edge length, even though it often works well for practical instances. Recent references on these topics can be found in the works of Edelsbrunner and Tan [101], Edelsbrunner, Tan, and Waupotitsch [102], and Dickerson, McElfresh, and Montague [83]. Rippa [195], and also Rippa and Schiff [194], gave other useful criteria in the context of approximating surfaces for which the Delaunay triangulation is optimal. Desnoguès [79] provides a good survey of polyhedral approximation.

An incremental algorithm that computes the Voronoi diagram of a set of points (see exercise 17.10) was given by Green and Sibson [113]. The Delaunay tree introduced by Boissonnat and Teillaud [31, 32] improves on the average performance when the sites are inserted in random order. This algorithm was made fully dynamic by Devillers, Meiser, and Teillaud [81]. An algorithm that proceeds by flipping diagonals was proposed by Lawson [144], then dynamized in the plane by Guibas, Knuth, and Sharir [117] who present a randomized analysis and also solutions to exercises 17.3 and 17.11. Its generalization to higher dimensions was studied by Joe [131, 132], Rajan [193], and Edelsbrunner and Shah [98], who provide a solution to exercise 17.13.

Lee proposed the first algorithm that computes Voronoi diagrams of higher orders in the plane [145]. He also gave a solution to exercise 17.14. The connection between Voronoi diagrams of order k and polytopes (exercise 17.15) was established by Aurenhammer [15]. Boissonnat, Devillers, and Teillaud [29] and also Mulmuley [174] proposed semi-dynamic or even fully dynamic algorithms that compute the Voronoi diagrams of all orders up to k in any dimension. Clarkson [67], Aurenhammer and Schwarzkopf [18], and also Agarwal, de Berg, and Matoušek [2] gave randomized algorithms that compute Voronoi diagrams of a single order k, rather than all the diagrams of orders $\leq k$.

The connection between Delaunay triangulations and Euclidean minimum spanning trees (see exercise 17.16) is discussed in the book by Preparata and Shamos [192], where one can also find a linear time algorithm that computes the Euclidean minimum spanning tree knowing the Delaunay triangulation. Conversely, Devillers [80] gave a randomized algorithm that computes the Delaunay triangulation of a set of n points in the plane knowing its Euclidean minimum spanning tree in expected time $O(n \log^* n)$.

For other references on Voronoi diagrams, the reader is referred to the book by Okabe, Boots, and Sugihara [182] or to the survey articles by Aurenhammer [16] and Fortune [105].