Chapter 18

Non-Euclidean metrics

In the previous chapter, we established a correspondence between the points in $\mathbb{E}^d$ and certain hyperplanes in $\mathbb{E}^{d+1}$, namely the hyperplanes tangent to the paraboloid $\mathcal{P}$. It is tempting to define the analogue of $\mathcal{V}(\mathcal{M})$ for a more general set of hyperplanes that may not necessarily be tangent to $\mathcal{P}$. In that case, the intersection of the n half-spaces lying above the hyperplanes is again a polytope whose proper faces, projected onto $\mathbb{E}^d$, form a cell complex that covers $\mathbb{E}^d$ entirely. This complex generalizes the Voronoi diagram and can be considered as the Voronoi diagram of a family of spheres, when the distance is defined as the power of a point with respect to one of the spheres. This interpretation, to be detailed in section 18.1, justifies the appellation *power diagrams* for such diagrams. These diagrams play a central role in several generalizations of Voronoi diagrams: in particular, we explore affine diagrams, which are Voronoi diagrams of point sites for a general quadratic distance (see section 18.2), and diagrams for weighted distances (see section 18.3).

Not all Voronoi diagrams for different metrics can be cast into power diagrams. For instance, polyhedral distances (and especially L_1 and L_∞) have important applications and are studied in section 18.4, and an application of hyperbolic Voronoi diagrams (see section 18.5) is given in the next chapter.

The representation of spheres introduced and used in the previous chapter is again very useful for computing power diagrams and hyperbolic Voronoi diagrams. In addition to this representation of spheres, we introduce in this chapter a new way of looking at Voronoi diagrams that is helpful for studying weighted diagrams, L_1 and L_∞ diagrams, and for the algorithms in the next chapter. Intuitively, the Voronoi diagram of a set $\mathcal{M}$ of points can be interpreted as the result of a growth process starting with the points in $\mathcal{M}$. Indeed, imagine crystals growing from each point of $\mathcal{M}$ at the same rate in all directions. The growth of a crystal stops at the points where it encounters another crystal, because of the constraint that the crystals may not interpenetrate. The crystal originating

at a point M_i in $\mathcal{M}$ covers the region that is reached by that crystal first, or in other words the points that are closer to M_i than to any other point in $\mathcal{M}$: this is exactly the Voronoi cell of M_i.

This growth process in $\mathbb{E}^d$ can be visualized in $\mathbb{E}^{d+1}$ by adding another coordinate, considered as the *time* elapsed since the start of the growth process. Thus $\mathbb{E}^d$ corresponds to the hyperplane $x_{d+1} = 0$ in $\mathbb{E}^{d+1}$, and the isotropic growth of a point M_i is a cone of revolution with vertex M_i and vertical (that is, parallel to the x_{d+1}-axis) axis. The faces of the Voronoi diagram appear as the projections onto $\mathbb{E}^d$ of faces on the lower envelope of the cones. If the sites do not start growing at the same time, the cones are translated vertically: this leads to Voronoi diagrams with additive weights. If the sites do not grow at the same rate, then the angles of the cones are different: the resulting Voronoi diagrams have multiplicative weights. If the sites do not grow isotropically (namely at the same rate in all directions), the cones are no longer cones of revolution: in this way we can generate Voronoi diagrams for the L_1 and L_∞ distances, and more generally for polyhedral distances (see exercise 19.3).

Throughout this chapter, the "distances" we consider are not exactly distance functions in the mathematical sense. In fact, we will only require that the distance function is increasing.

18.1 Power diagrams

18.1.1 Definition and computation

Let $\mathcal{S} = \{\Sigma_1, \ldots, \Sigma_n\}$ be a set of n spheres in $\mathbb{E}^d$. To each Σ_i corresponds a region $P(\Sigma_i)$ of $\mathbb{E}^d$, consisting of the points whose power with respect to Σ_i is smaller than their powers with respect to the other spheres:

$$P(\Sigma_i) = \{X \in \mathbb{E}^d \ : \ \forall j \neq i, \ \Sigma_i(X) \leq \Sigma_j(X)\}.$$

The region $P(\Sigma_i)$ is the intersection of a finite number of half-spaces (bounded by the radical hyperplanes H_{ij}, $j = 1, \ldots, n$, $j \neq i$). It is therefore a convex polytope, occasionally empty or unbounded. The $P(\Sigma_i)$'s and their faces form a cell complex which covers $\mathbb{E}^d$: this complex is called the *power diagram* of $\mathcal{S}$ and we denote it by $\mathcal{P}ow(\mathcal{S})$ (see figure 18.1).

As in the previous chapter, we map a sphere Σ in $\mathbb{E}^d$, centered at C and of equation $\Sigma(X) = 0$, to the point $\phi(\Sigma) = (C, \Sigma(0))$ in $\mathbb{E}^{d+1}$. The hyperplane polar to $\phi(\Sigma)$ with respect to the paraboloid $\mathcal{P}$ is denoted by $\phi(\Sigma)^*$: if $\mathbb{E}^d$ is embedded in $\mathbb{E}^{d+1}$ as the hyperplane $x_{d+1} = 0$, then $\phi(\Sigma)^*$ is the hyperplane that intersects $\mathcal{P}$ along the quadric obtained by lifting Σ onto $\mathcal{P}$ (see figure 17.4).

Let $\mathcal{P}(\mathcal{S})$ be the intersection of the half-spaces bounded below by the polar hyperplanes $\phi(\Sigma_1)^*, \ldots, \phi(\Sigma_n)^*$ to the points $\phi(\Sigma_1), \ldots, \phi(\Sigma_n)$.

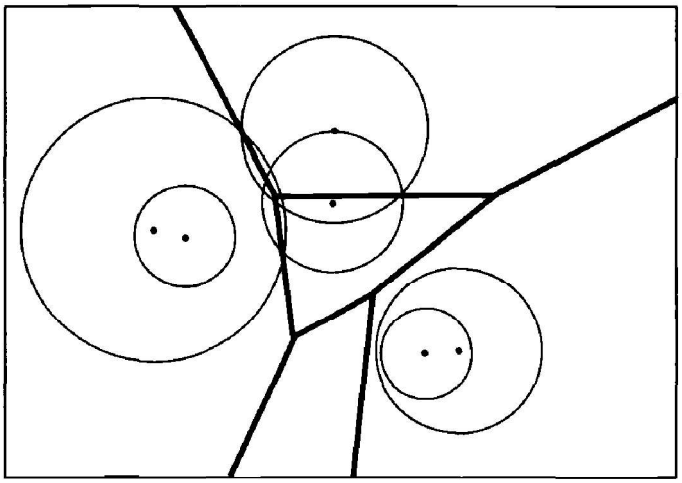

Figure 18.1. A power diagram.

Theorem 18.1.1 *The power diagram $\mathcal{P}ow(\mathcal{S})$ of a set $\mathcal{S} = \{\Sigma_1, \ldots, \Sigma_n\}$ of n spheres in $\mathbb{E}^d$ is a cell complex in $\mathbb{E}^{d+1}$. Its faces are obtained by projecting the proper faces of the unbounded $(d+1)$-polytope $\mathcal{P}(\mathcal{S})$, the intersection of n half-spaces in $\mathbb{E}^{d+1}$ bounded below by the polar hyperplanes $\phi(\Sigma_1)^*, \ldots, \phi(\Sigma_n)^*$.*

Proof. Let $\underline{A}$ be a point on a facet of $\mathcal{P}(\mathcal{S})$ which is supported by the polar hyperplane $\phi(\Sigma_i)^*$, and let A be its projection onto $\mathbb{E}^d$. The power of A with respect to Σ_i is the signed vertical distance from $\underline{A}$ to $\phi(A)$ (lemma 17.2.3). Since $\underline{A}$ belongs to a facet of $\mathcal{P}(\mathcal{S})$, the power of A with respect to Σ_i is less than or equal to the power of A with respect to any other sphere in $\mathcal{S}$. In other words, A belongs to the cell $P(\Sigma_i)$ that corresponds to Σ_i in the power diagram. □

Note that when the hyperplanes $\phi(\Sigma_i)^*$ are in general position, $\mathcal{P}(\mathcal{S})$ is a simple polytope in $\mathbb{E}^{d+1}$, so each vertex is incident to $d+1$ hyperplanes. In terms of the spheres Σ_i, this general position assumption means that no subset of $d+2$ spheres in $\mathcal{S}$ are orthogonal to a common sphere in $\mathbb{E}^d$, or equivalently that no point in $\mathbb{E}^d$ has the same power with respect to $d+2$ spheres in $\mathcal{S}$. In this case, we say that the spheres are in *general position*. The power diagram $\mathcal{P}ow(\mathcal{S})$ is a cell complex whose vertices have the same power with respect to $d+1$ spheres in $\mathcal{S}$ (and are therefore the centers of spheres orthogonal to $d+1$ spheres in $\mathcal{S}$), and have a greater power with respect to the other spheres in $\mathcal{S}$. More generally, a k-face in $\mathcal{P}ow(\mathcal{S})$ is formed by the points that have the same power with respect to $d+1-k$ given spheres in $\mathcal{S}$, and a greater power with respect to the other spheres in $\mathcal{S}$.

Corollary 18.1.2 *The complexity of the power diagram of n spheres in $\mathbb{E}^d$ is $\Theta(n^{\lceil d/2 \rceil})$. The diagram can be computed in time $O(n \log n + n^{\lceil d/2 \rceil})$, which is*

optimal in the worst case.

Remark 1. In the case of Voronoi diagrams, all the hyperplanes $\phi(M_i)^*$ are tangent to the paraboloid $\mathcal{P}$, and each of them contributes a facet to $\mathcal{V}(\mathcal{M})$. For power diagrams, however, a polar hyperplane $\phi(\Sigma_i)^*$ does not necessarily contribute a face to $\mathcal{P}(\mathcal{S})$. Such a hyperplane is called *redundant.* In the power diagram, it means that the cell $P(\Sigma_i)$ is empty: Σ_i does not contribute a cell to $\mathcal{P}ow(\mathcal{S})$.

Remark 2. There is no particular difficulty if some, even all, the spheres in $\mathcal{S}$, are imaginary. This fact is used in section 18.3.2.

Remark 3. Any polytope in $\mathbb{E}^{d+1}$ that is the intersection of upper half-spaces corresponds to a power diagram: if $H_1, \ldots, H_n$ are the hyperplanes that bound these half-spaces, then their upper envelope projects onto the power diagram of the spheres $\phi^{-1}(H_1^*), \ldots, \phi^{-1}(H_n^*)$.

18.1.2 Higher-order power diagrams

As was done for Voronoi diagrams in section 17.4, we may define power diagrams of higher orders.

Let $\mathcal{S}_k$ be a subset of $\mathcal{S}$ of size k. We call the power cell of $\mathcal{S}_k$ the set $P(\mathcal{S}_k)$ of points in $\mathbb{E}^d$ that have a smaller power with respect to any sphere in $\mathcal{S}_k$ than to any sphere in $\mathcal{S} \setminus \mathcal{S}_k$:

$$P(\mathcal{S}_k) = \{X \in \mathbb{E}^d \ : \ \forall \Sigma_i \in \mathcal{S}_k, \ \forall \Sigma_j \in \mathcal{S} \setminus \mathcal{S}_k, \ \Sigma_i(X) \leq \Sigma_j(X)\}.$$

Consider all the subsets of size k of $\mathcal{S}$ whose corresponding power cell is not empty. These regions and their faces form a cell complex that covers $\mathbb{E}^d$ entirely, and that is called the *power diagram of order k* of $\mathcal{S}$. We denote it by $\mathcal{P}ow_k(\mathcal{S})$.

This fact is a consequence of the theorem below, whose proof closely resembles that of theorem 17.4.1. This theorem clarifies the links between power diagrams of order k in $\mathbb{E}^d$ and faces at level k in the arrangement of n hyperplanes in $\mathbb{E}^{d+1}$. As usual, the Euclidean space of dimension d is identified with the hyperplane $x_{d+1} = 0$ in the space $\mathbb{E}^{d+1}$ of dimension $d+1$, and $\phi(\Sigma)^*$ stands for the polar hyperplane of $\phi(\Sigma)$.

Theorem 18.1.3 *Consider a set $\mathcal{S} = \{\Sigma_1, \ldots, \Sigma_n\}$ of spheres in $\mathbb{E}^d$, and let $\mathcal{A}$ be the arrangement of their polar hyperplanes $\phi(\Sigma_1)^*, \ldots, \phi(\Sigma_n)^*$. The power diagram of order k, $\mathcal{P}ow_k(\mathcal{S})$, is a cell d-complex in $\mathbb{E}^d$. Its cells are the vertical projections of the cells at level k in the arrangement $\mathcal{A}$, the reference point being on the x_{d+1}-axis above all the hyperplanes $\phi(\Sigma_i)^*$, $i = 1, \ldots, n$. The l-faces of $\mathcal{P}ow_k(\mathcal{S})$, $l < d$, are obtained by projecting the l-faces common to at least two cells of $\mathcal{A}$ at level k.*

From theorems 14.5.1 and 14.5.3, we derive the following result.

Theorem 18.1.4 *The complexity of the first k power diagrams of a set of n spheres in $\mathbb{E}^d$ is $O(n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil})$. These k diagrams can be computed in time $O(n^{\lfloor (d+1)/2 \rfloor} k^{\lceil (d+1)/2 \rceil})$ if $d \geq 3$, and in time $O(nk^2 \log \frac{n}{k})$ if $d = 2$.*

18.2 Affine diagrams

The notion of a Voronoi diagram can be extended to more general sites or to non-Euclidean distances. A particularly interesting extension occurs when the locus of points equidistant from two sites is a hyperplane: in this case, the diagram is called an *affine diagram.* Voronoi diagrams and power diagrams are affine diagrams, and we will show that any affine diagram is a power diagram. Moreover, certain non-affine diagrams can be derived from an affine diagram and therefore from a power diagram: this is notably the case of diagrams with additive or multiplicative weights studied in section 18.3.

18.2.1 Affine diagrams and power diagrams

An affine diagram is a diagram defined for object sites and for a distance such that the set of points equidistant from two objects is a hyperplane. The cells of such diagrams are thus convex polytopes and affine diagrams can be identified with cell complexes.

To any affine diagram of n objects corresponds a set of $\binom{n}{2}$ perpendicular bisectors H_{ij}, $1 \leq i < j \leq n$. These hyperplanes must satisfy the relations

$$H_{ij} \cap H_{jk} = H_{ij} \cap H_{ik} = H_{ik} \cap H_{jk} \stackrel{\text{def}}{=} I_{ijk}$$

for any $1 \leq i < j < k \leq n$.

We say that the diagram is *simple* if the I_{ijk} are disjoint and not empty.

Theorem 18.2.1 *Any simple affine diagram in $\mathbb{E}^d$ is the power diagram of a set of spheres in $\mathbb{E}^d$.*

Proof. We embed $\mathbb{E}^d$ in $\mathbb{E}^{d+1}$ as the hyperplane $x_{d+1} = 0$. The proof consists of constructing a set of n hyperplanes $P_1, \ldots, P_n$ in $\mathbb{E}^{d+1}$ such that the vertical projection of $P_i \cap P_j$ for $i < j$ is exactly H_{ij}. Assuming these hyperplanes are known, to each P_i corresponds a sphere $\Sigma_i = \phi^{-1}(P_i^*)$ whose polar hyperplane is exactly P_i: Σ_i is also the projection on $\mathbb{E}^d$ of the intersection of P_i with the paraboloid $\mathcal{P}$. Hence H_{ij} is the radical hyperplane of Σ_i and Σ_j for all i and j

such that $1 \leq i < j \leq n$. It follows that the affine diagram is exactly the power diagram of the spheres Σ_i, $i = 1, \ldots, n$.

We now show how to build the P_i's. Denote by h_{ij} the vertical projection of H_{ij} onto P_i (note that $i < j$ by the definition of H_{ij}).

Let us take for P_1 any non-vertical hyperplane, and for P_2 any non-vertical hyperplane that intersects P_1 along h_{12}. For $k \geq 3$, we must take for P_k the hyperplane that intersects P_1 along h_{1k} and P_2 along h_{2k}: such a hyperplane exists because h_{1k} and h_{2k} intersect along the affine subspace of dimension $d-2$ that is the projection of I_{12k} onto P_1, P_2, or P_k.

It remains to see that the vertical projection of $P_i \cap P_j$ is exactly H_{ij}. By construction, this is true for $P_1 \cap P_2$, $P_1 \cap P_j$, and $P_2 \cap P_j$, $j \geq 3$. For $3 \leq i < j \leq n$, we know that $P_i \cap P_j \cap P_1$ projects onto $\mathbb{E}^d$ along I_{1ij}, and that $P_i \cap P_j \cap P_2$ projects onto $\mathbb{E}^d$ along I_{2ij}. The diagram being simple, I_{1ij} and I_{2ij} must be distinct. The projection of $P_i \cap P_j$ must therefore contain I_{1ij} and I_{2ij}, and hence also their affine hull which is nothing other than H_{ij}. □

Below, we rather use

Theorem 18.2.2 *The affine diagram whose hyperplanes H_{ij} have equations*

$$-2(C_i - C_j) \cdot X + \sigma_i - \sigma_j = 0$$

is the power diagram of the spheres Σ_i, $i = 1, \ldots, n$ centered at C_i and with respect to which the origin has power σ_i.

Proof. We may simply check that the equation of H_{ij} can be written as $\Sigma_i(X) - \Sigma_j(X) = 0$, which is exactly that of the radical hyperplane of Σ_i and Σ_j (see subsection 17.2.6). □

18.2.2 Diagrams for a general quadratic distance

Consider two points X and A in $\mathbb{E}^d$. By the *general quadratic distance* from A to X, we mean the quantity

$$\delta_Q(X, A) = (X - A)\Delta(X - A)^t + p(A),$$

where Δ is a real symmetric $d \times d$ matrix, and where $p(A)$ is a real number.

The diagrams encountered so far are all particular cases of Voronoi diagrams for a quadratic distance:

- Standard Voronoi diagrams are obtained for $\Delta = \mathbb{I}_d$ and $p(X) = 0$.
- Furthest-point Voronoi diagrams (introduced as diagrams of n points of order $n-1$) are obtained for $\Delta = -\mathbb{I}_d$ and $p(X) = 0$.

- Power diagrams correspond to $\Delta = \mathbb{I}_d$ and $p(X) \neq 0$.

For any pair of points A and B, the set of points X that are equidistant from A and B is the hyperplane H_{AB} equation

$$H_{AB} \; : \; 2(B-A)\Delta X^t + A\Delta A^t - B\Delta B^t + p(A) - p(B) = 0.$$

The Voronoi diagram of a finite set of points for a general quadratic distance is thus an affine diagram by theorem 18.2.2.

Theorem 18.2.3 *The Voronoi diagram of n points for an arbitrary general quadratic distance in $\mathbb{E}^d$ has complexity $\Theta(n^{\lceil d/2 \rceil})$. It can be computed in time $\Theta(n \log n + n^{\lceil d/2 \rceil})$.*

18.3 Weighted diagrams

This section introduces two kinds of diagrams which are not affine. They are defined for finite sets of point sites and for a Euclidean distance that is weighted additively or multiplicatively. Each distance is appropriately defined in the subsection below.

These diagrams are not cell complexes those we have been studying so far. Nevertheless, they can be given a facial structure that is similar to that of cell complexes. Consider the equivalence relation shared by the points in $\mathbb{E}^d$ that have the same nearest neighbors. The equivalence classes subdivide $\mathbb{E}^d$ into (open) regions whose closures we call the *faces* of the diagram. The cells of the diagram span $\mathbb{E}^d$ entirely and the intersection of two faces is a (possibly empty) collection of lower-dimensional faces.

As we see below, the faces of these diagrams are not polytopes, and may not even be connected. Nevertheless, these weighted diagrams can be derived simply from power diagrams.

18.3.1 Weighted diagrams with additive weights

Let $\mathcal{M} = \{M_1, \ldots, M_n\}$ be a set of n points in $\mathbb{E}^d$. To each M_i corresponds a real r_i called the *weight* of M_i. The *additive weighted distance*, or *additive distance* for short, from a point X in $\mathbb{E}^d$ to M_i is the quantity

$$\delta_+(X, M_i) = \|XM_i\| - r_i.$$

The *diagram of $\mathcal{M}$ with additive weights* is defined like the Voronoi diagram except that the distance used is not the Euclidean distance but the additive distance defined above. This diagram is denoted by $\mathcal{V}or_+(\mathcal{M})$. An instance is

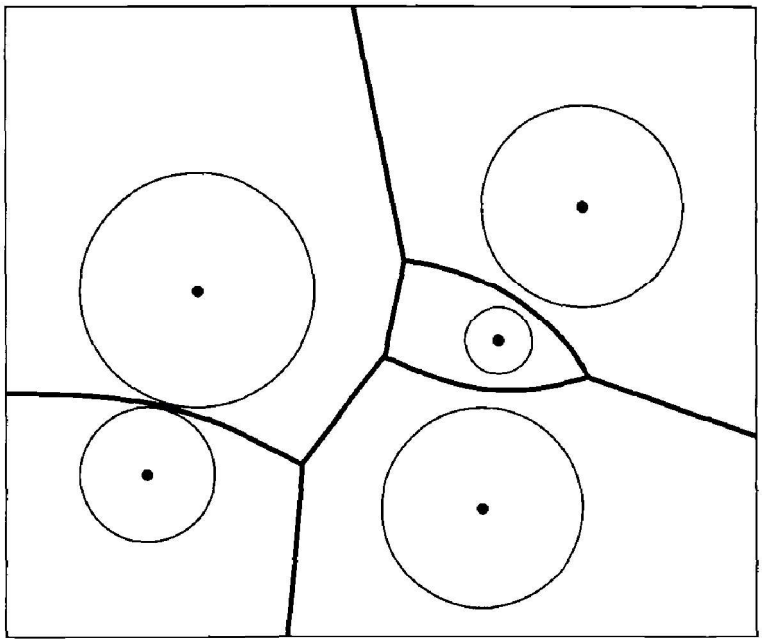

Figure 18.2. A diagram with additive weights. Sites are the centers and their corresponding weights are the radii of the circles. In this example, the diagram of the points with additive weights is also the Voronoi diagram of the circles for the Euclidean metric.

shown in figure 18.2. We observe that it is the first example of a non-affine diagram shown in this book.

Note that adding the same constant to all the points does not modify the diagram. This lets us assume that all the r_i's are non-negative.

The representation of spheres introduced in section 17.2 is not very helpful here and we use another which shows a very natural correspondence between weighted Voronoi diagrams and affine diagrams in dimension $d+1$.

Consider the sphere Σ_i in $\mathbb{E}^d$ centered at M_i with radius r_i, and let ψ be the bijection that maps Σ_i to the point $\psi(\Sigma_i) = (M_i, r_i) \in \mathbb{E}^{d+1}$.

The spheres of zero radius correspond to the points in the hyperplane of equation $x_{d+1} = 0$ in $\mathbb{E}^{d+1}$.

The points at additive distance r from M_i can be considered as the centers of the spheres of radius $|r|$ tangent to Σ_i, that are inside or outside Σ_i according to whether r is negative or positive. The images under ψ of these spheres generate a cone of revolution $\mathcal{C}(\Sigma)$ of equation

$$\mathcal{C}(\Sigma) \; : \; x_{d+1} = \|XC\| - r$$

which has apex $(C, -r)$, is symmetrical to $\psi(\Sigma)$ with respect to the hyperplane $x_{d+1} = 0$, and has an aperture angle of $\frac{\pi}{4}$. The vertical projection I_X of a point

X in $\mathbb{E}^d$ on the cone $C(\Sigma)$ is the image under ψ of the sphere centered at X and tangent to Σ. The signed vertical distance from X to I_X equals the additive distance from X to C weighted by r.

To each sphere Σ_i, $i = 1, \ldots, n$, corresponds the cone $C(\Sigma_i)$, also denoted by $\mathcal{C}_i$. It follows from the discussion above that the projection of the lower envelope of the cones $\mathcal{C}_i$ onto $\mathbb{E}^d$ is exactly $\mathcal{V}or_+(\mathcal{M})$.

The set of points in $\mathbb{E}^d$ that are equidistant (with respect to the additive distance) from two points of $\mathcal{M}$ is thus the projection of the intersection of two cones. This intersection is a quadric contained in a hyperplane. Indeed, we have

$$\begin{aligned} \mathcal{C}_1 \; : \; (x_{d+1} + r_1)^2 &= XM_1^2, \;\; x_{d+1} + r_1 > 0, \\ \mathcal{C}_2 \; : \; (x_{d+1} + r_2)^2 &= XM_2^2, \;\; x_{d+1} + r_2 > 0. \end{aligned}$$

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

$$H_{12} \; : \; -2(M_1 - M_2) \cdot X - 2(r_1 - r_2)x_{d+1} + M_1^2 - r_1^2 - M_2^2 + r_2^2 = 0.$$

This and theorem 18.2.2 show that there exists a correspondence between the diagram $\mathcal{V}or_+(\mathcal{M})$ and the power diagram of the spheres Σ_i' in $\mathbb{E}^{d+1}$ $(i = 1, \ldots, n)$, where Σ_i' is centered at $\psi(\Sigma_i)$ and has radius $r_i\sqrt{2}$ (see figure 18.3). More precisely, the cell of $\mathcal{V}or_+(\mathcal{M})$ that corresponds to M_i is the projection of the intersection of the cone $\mathcal{C}_i$ with the cell of the power diagram corresponding to the sphere Σ_i'. Indeed, X is in $\mathcal{V}or_+(M_i)$ 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 cones $\mathcal{C}_j$, $j \neq i$. In other words, the coordinates (X, x_{d+1}) of X_i must obey

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

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

The additive diagram can be computed using the following algorithm:

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 onto $\mathbb{E}^d$ the intersection with the cone $\mathcal{C}_i$ of the cell of the power diagram that corresponds to Σ_i'.

The power diagram of the Σ_i' can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$. 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 $O(n^{\lfloor d/2 \rfloor + 1})$. We have thus proved that:

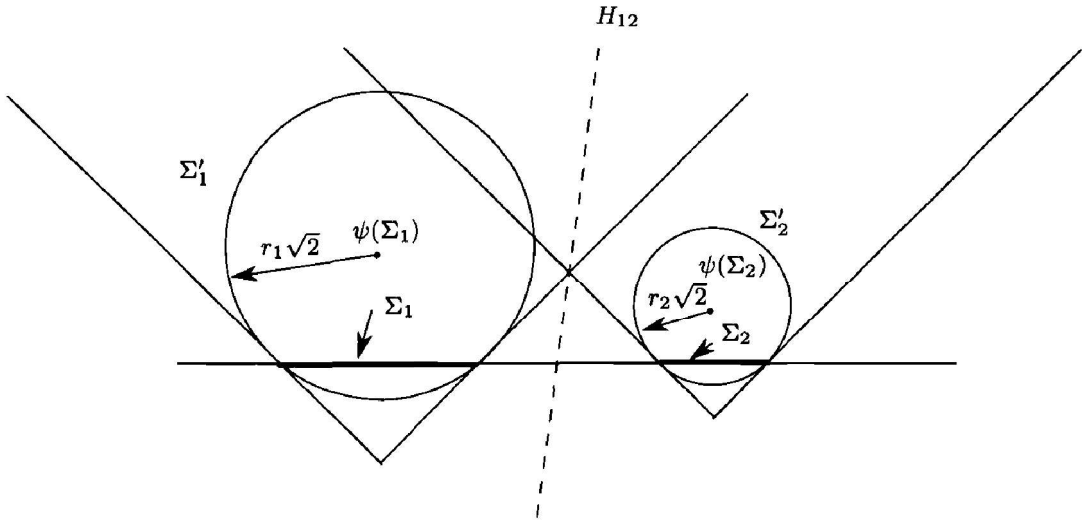

Figure 18.3. Any Voronoi diagram for the additive distance can be derived from a power diagram in $\mathbb{E}^{d+1}$.

Theorem 18.3.1 *The Voronoi diagram of a set of n points in $\mathbb{E}^d$ with additive weights has complexity $O(n^{\lfloor d/2 \rfloor+1})$ and can be computed in time $O(n^{\lfloor d/2 \rfloor+1})$.*

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, however, as we now show. We also conjecture that it is not optimal in any even dimension.

In the plane, we have seen that additive diagrams can be thought of as the projection onto $\mathbb{E}^2$ of the lower envelope of cones with vertical axis and aperture angle $\frac{\pi}{4}$. Therefore, each cell is connected. Moreover, the vertices of the diagram are incident to exactly three edges, under the general position assumption, and these edges are arcs of hyperbolas, each of which is the projection of the intersection of two cones. Euler's relation shows that the diagram has complexity $O(n)$. A perturbation argument shows that the general position assumption is not restrictive, since allowing degeneracies only merges some vertices and makes some edges disappear. In section 19.1, it is shown that such a diagram can be computed in optimal time $O(n \log n)$.

18.3.2 Weighted diagrams with multiplicative weights

Let $\mathcal{M} = \{M_1, \ldots, M_n\}$ be a set of n point sites in $\mathbb{E}^d$. To each M_i corresponds a positive real number $p(M_i)$ called the *weight* of M_i. To simplify the presentation, we suppose that the $p(M_i)$'s are all distinct, but the extension to the more general case presents no additional difficulties.

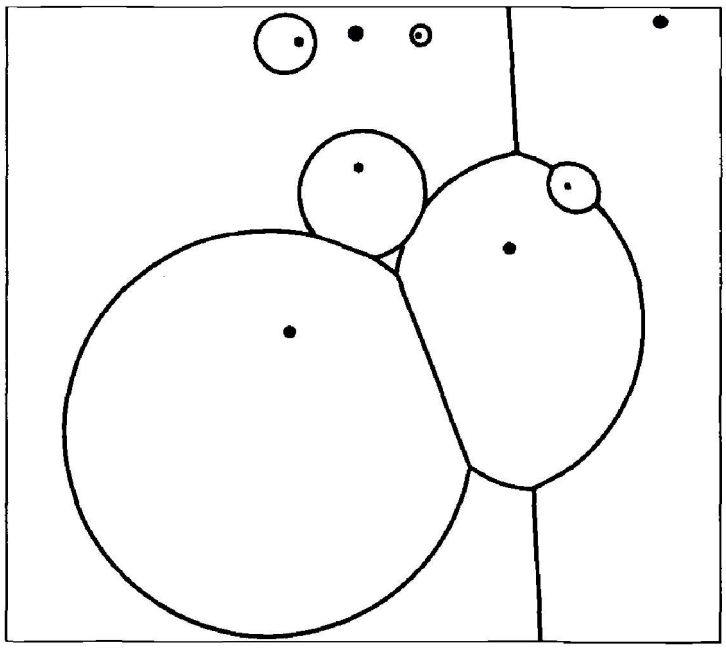

Figure 18.4. A diagram with multiplicative weights. Sites are represented by small disks, and the weight of a site is inversely proportional to the diameter of the disk.

By the *distance with multiplicative weights*, or *multiplicative distance* for short, from a point X to a point M_i, we mean the quantity

$$\delta_*(X, M_i) = p(M_i)\,\|XM_i\|.$$

The *Voronoi diagram of $\mathcal{M}$ for the multiplicative distance* is defined like the Voronoi diagram, except that the distance is not the Euclidean distance but rather the multiplicative distance. We denote this diagram by $\mathcal{V}or_*(\mathcal{M})$ (see figure 18.4). Observe that a cell of the diagram need not be connected.

The set of points at equal multiplicative distance from two sites M_i and M_j is a sphere Σ_{ij} of equation

$$p_i\,(X - M_i)^2 = p_j\,(X - M_j)^2$$

with $p_i = p(M_i)^2$. In normalized form, we obtain

$$X^2 - 2\frac{p_iM_i - p_jM_j}{p_i - p_j} \cdot X + \frac{p_iM_i^2 - p_jM_j^2}{p_i - p_j} = 0.$$

Using the representation of section 17.2, this sphere is represented in $\mathbb{E}^{d+1}$ as the point

$$\phi(\Sigma_{ij}) = \left(\frac{p_iM_i - p_jM_j}{p_i - p_j}, \frac{p_iM_i^2 - p_jM_j^2}{p_i - p_j}\right).$$

Its polar hyperplane H_{ij} with respect to the paraboloid $\mathcal{P}$ has equation

$$H_{ij}(X, x_{d+1}) = (p_i - p_j)x_{d+1} - 2p_iM_i \cdot X + 2p_jM_j \cdot X + p_iM_i^2 - p_jM_j^2 = 0.$$

The hyperplanes H_{ij} are the radical hyperplanes of spheres Σ_i in $\mathbb{E}^{d+1}$ ($i = 1, \ldots, n$). The sphere Σ_i, possibly imaginary, is centered at $(p_iM_i, -\frac{p_i}{2})$, and with respect to it the origin has power $p_iM_i^2$. This establishes a correspondence between the diagram $\mathcal{V}or_*(\mathcal{M})$ and the power diagram of the spheres Σ_i. More precisely, the cell $V_*(M_i)$ in $\mathcal{V}or_*(\mathcal{M})$ that corresponds to M_i is the projection of the intersection of the paraboloid $\mathcal{P}$ with the cell $P(\Sigma_i)$ that corresponds to Σ_i in the power diagram of the Σ_i's. Indeed, if X is a point in $\mathbb{E}^d$ and $\phi(X)$ is its vertical projection onto the paraboloid $\mathcal{P}$ of equation $x_{d+1} = X^2$, then we have

$$\begin{aligned} X \in V_*(M_i) &\iff p_i(X - M_i)^2 \leq p_j(X - M_j)^2 \qquad \forall j \neq i \\ &\iff H_{ij}(X, X^2) \leq 0 \qquad \forall j \neq i \\ &\iff \Sigma_i(\phi(X)) \leq \Sigma_j(\phi(X)) \qquad \forall j \neq i \\ &\iff \phi(X) \in P(\Sigma_i). \end{aligned}$$

An algorithm that computes the diagram of $\mathcal{M}$ with multiplicative weights is:

1. Compute Σ_i, for $i = 1, \ldots, n$.
2. Compute the power diagram of the Σ_i's.
3. For $i = 1, \ldots, n$, project the intersection of the cell that corresponds to Σ_i in the power diagrams of the Σ_i's with the paraboloid $\mathcal{P}$.

This proves the following theorem.

Theorem 18.3.2 *The Voronoi diagram of a set of n points in $\mathbb{E}^d$ with multiplicative weights has complexity $O(n^{\lfloor d/2 \rfloor + 1})$ and can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$.*

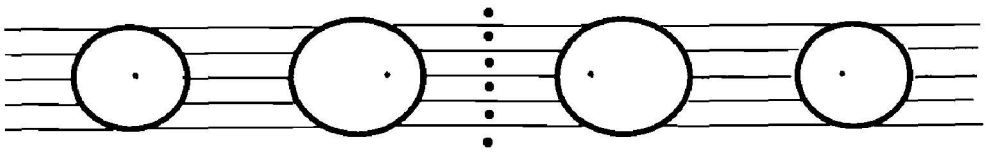

Figure 18.5. An instance of a quadratic multiplicative diagram in dimension 2: $\frac{n}{2}$ points are put on a given vertical line and are given the same weight, while $\frac{n}{2}$ other points are aligned on a horizontal line and have the same weight, which is much larger than the weight given to the points in the first half.

This result is optimal in odd dimensions, since in that case these bounds match those of the Voronoi diagram of n points in $\mathbb{E}^d$ for the Euclidean distance. It is also optimal in even dimensions (see exercise 18.4). Figure 18.5 shows a quadratic multiplicative diagram in dimension 2.

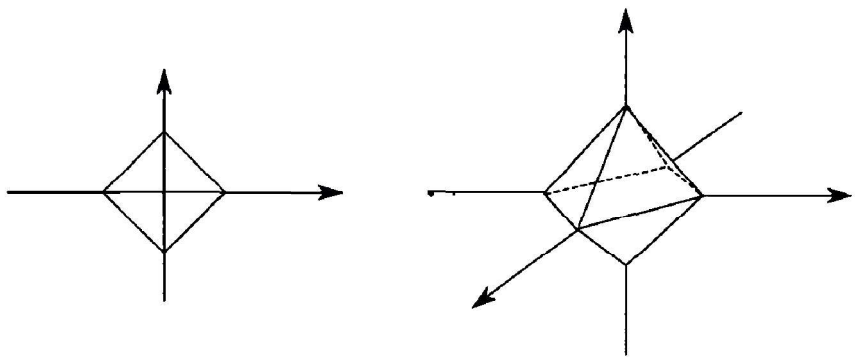

Figure 18.6. Co-cubes in dimensions 2 and 3.

18.4 L_1 and L_∞ metrics

The L_1 distance of a point $X = (x_1, \ldots, x_d)$ in $\mathbb{E}^d$ to a point $M = (m_1, \ldots, m_d)$ in $\mathbb{E}^d$ is defined as

$$\delta_1(X, M) = \sum_{i=1}^{d} |x_i - m_i|.$$

The points at a given distance r from M are thus on a polytope whose vertices are given by their coordinates $x_i = m_i \pm r$ and $x_j = m_j$ if $i \neq j$, for $j = 1, \ldots, d$. In dimension 2 this polytope is a tilted square, and in dimension 3 it is a regular octahedron (see figure 18.6). This polytope is dual to the cube and we call it a *co-cube*. Henceforth, a co-cube always means a polytope dual to a cube whose edges are parallel coordinate axes.

Let $\mathcal{M} = \{M_1, \ldots, M_n\}$ be a set of n point sites in $\mathbb{E}^d$. The *Voronoi diagram of $\mathcal{M}$ for the L_1 distance* is defined similarly to the Voronoi diagram, except that the distance used in the definition of the cells is not the Euclidean distance but the L_1 distance. It is denoted by $\mathcal{V}or_{L_1}(\mathcal{M})$ (see figure 18.7).

We can define a facial structure for this diagram by using the equivalence relation R shared by the points in $\mathbb{E}^d$ that have the same subset of nearest neighbors. The equivalence classes of R subdivide the space $\mathbb{E}^d$ in open regions whose closures are called the *faces* of the diagram. The faces of the diagram are piecewise affine.

If the points in $\mathbb{E}^d$ are identified with the hyperplane $x_{d+1} = 0$ in $\mathbb{E}^{d+1}$, then, in a way similar to what was explained in subsection 18.3.1, to each point M_i there corresponds a pyramid $\mathcal{P}_i$ of equation

$$x_{d+1} = \delta_1(X, M_i).$$

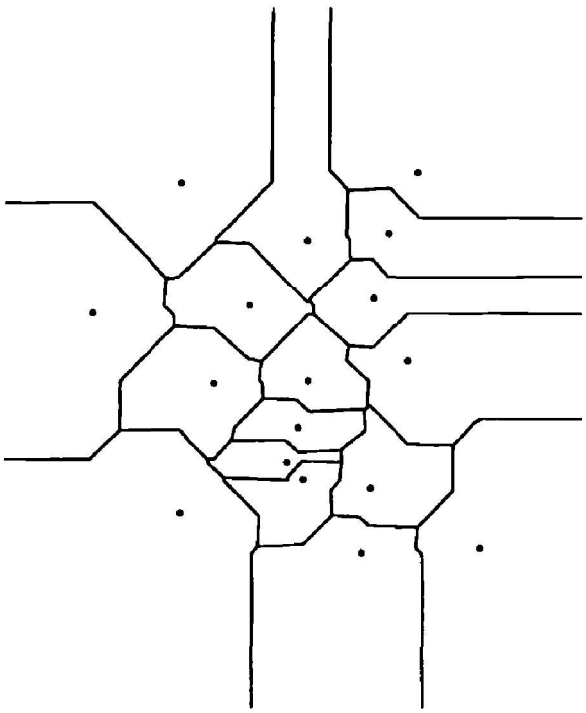

Figure 18.7. Diagram for the L_1 metric.

Let us consider the lower envelope of the $\mathcal{P}_i$'s, that is, the graph of the function $\min_{1 \leq i \leq n} \delta_1(X, M_i)$. The portion of the lower envelope that belongs to any $\mathcal{P}_i$ projects onto the hyperplane $x_{d+1} = 0$ as the cell of the diagram $\mathcal{V}or_{L_1}(\mathcal{M})$ that corresponds to M_i. The facets of the $\mathcal{P}_i$'s form a collection of d-pyramids. The lower envelope of these pyramids is a collection of d-faces, and their lower-dimensional faces include all the lower-dimensional faces of the lower envelope of the $\mathcal{P}_i$'s. The vertical projections onto $x_{d+1} = 0$ of the d-faces of the lower envelope of the pyramids form a refinement of the faces of the diagram $\mathcal{V}or_{L_1}(\mathcal{M})$. The complexity of the diagram $\mathcal{V}or_{L_1}(\mathcal{M})$ can thus be bounded by combining theorem 16.3.2 and exercise 16.1, which bound the complexity of the lower envelope of n d-simplices in $\mathbb{E}^{d+1}$. This yields

$$|\mathcal{V}or_{L_1}(\mathcal{M})| = O(n^d \alpha(n)).$$

This bound is almost tight for certain sets of points that are not in general position (see exercise 18.9). We conjecture, however, that for points in general enough position, this bound is not attained and that these diagrams have the same complexity as their Euclidean counterparts. Later on, we show that this is indeed the case in dimension 2, for which we give a linear bound. It is also the case in dimension 3 (see exercise 18.10). If $d = 2$, the bisector for the L_1-distance of two points is, in general, a polygonal line formed by three linear pieces; if the line connecting the two points is parallel to one of the main bisectors, however,

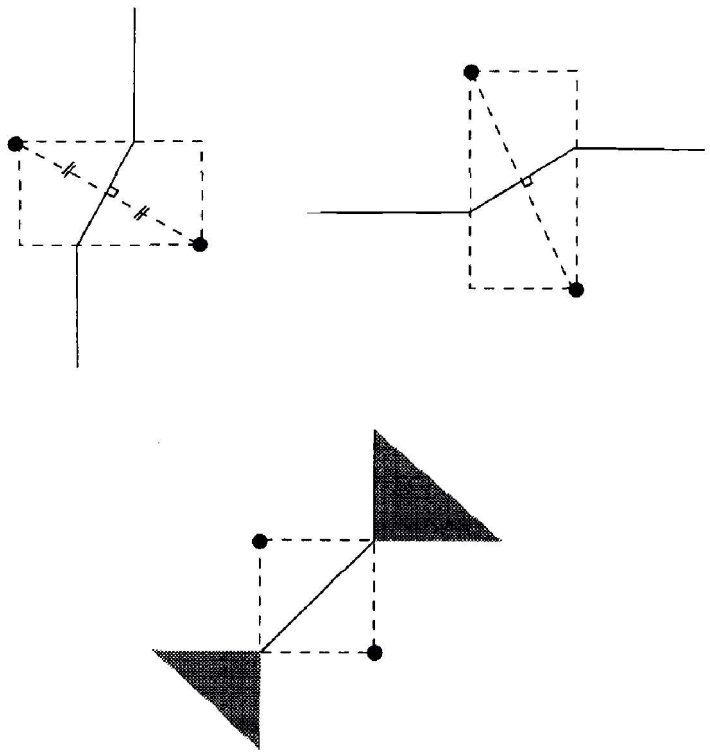

Figure 18.8. Bisectors for the L_1 distance. If the line connecting the two points is parallel to one of the main bisectors, the bisector is not a polygonal line.

the L_1-bisector is no longer a polygonal line and contains two faces of dimension 2 (see figure 18.8).

We say that two points are in L_1-*general position* if no two points are connected by a line parallel to one of the main bisectors, and if no four points belong to a common co-cube. In this case, the bisectors are polygonal lines formed of three line segments, and $\mathcal{V}or_{L_1}(\mathcal{M})$ contains n connected cells: indeed, for any $i \in \{1, \ldots, n\}$, the cell $V_1(M_i)$ that corresponds to a point M_i is star-shaped with respect to M_i (meaning that if $X \in V_1(M_i)$, then the segment XM_i is contained in $V_1(M_i)$), and is therefore connected. Moreover, each vertex in the diagram is incident to two or three edges because of the L_1-general position assumption. The diagram $\mathcal{V}or_{L_1}(\mathcal{M})$ is therefore a planar map with n cells whose vertices have degree two or three and whose edges consist of at most three segments. Euler's relation then shows that the complexity of the diagram is $O(n)$.

If the points are not in L_1-general position, then some regions may correspond to pairs of points (see figure 18.8) and some vertices may be of degree higher than 3. This second complication can be straightened out by simply perturbing the diagram so as to replace each vertex of degree $k > 3$ by a small polygonal chain with $k - 2$ vertices of degree 3 and $k - 3$ edges. The number of faces does not

increase in the process, and the number of vertices increases by the same amount as the number of edges; hence Euler's relation still guarantees that the complexity of the diagram is $O(n)$. The first complication, however, is more serious and may allow the size to grow up to quadratic: exercise 18.9 presents such an example and a way to avoid this problem. The example generalizes to higher dimensions and the lower bound $\Omega(n^d)$ may be shown to hold for the complexity of Voronoi diagrams of n points in $\mathbb{E}^d$ for the L_1 distance.

If the points are in L_1-general position, the complexity of the diagram is thus $O(n)$ in dimension 2 and the algorithm that computes the lower envelope of n triangles in space (see subsection 16.3.3) can be used to compute this diagram in time $O(n \log^2 n)$ (see corollary 16.3.3). An optimal algorithm exists that computes such a diagram in time $O(n \log n)$ (see exercise 19.2).

The situation for the L_∞ distance is very similar to the one just described for the L_1 distance. Its complexity in dimensions higher than 3 is easier to analyze, however. The L_∞ *distance* of a point $X = (x_1, \ldots, x_d)$ in $\mathbb{E}^d$ from a point $M = (m_1, \ldots, m_d)$ in $\mathbb{E}^d$ is given by

$$\delta_\infty(X, M) = \max_{i=1,\ldots,d} |x_i - m_i|.$$

The points at a distance r from M are thus on a cube centered at M whose facets are parallel to the coordinate axes, and whose side is $2r$.

The *Voronoi diagram of $\mathcal{M}$ for the L_∞ distance* is denoted by $\mathcal{V}or_{L_\infty}(\mathcal{M})$. An instance is shown in figure 18.9.

The cells of this diagram can be obtained by projecting onto the hyperplane $x_{d+1} = 0$ in $\mathbb{E}^{d+1}$ the cells on the lower envelope of the n pyramids $\mathcal{Q}_i$ of equation

$$x_{d+1} = \delta_\infty(X, M_i).$$

The facets of the $\mathcal{Q}_i$'s form a collection of d-pyramids. The faces on the lower envelope of these pyramids form a refinement of the faces on the lower envelope of the $\mathcal{Q}_i$'s. Hence, the vertical projections onto the hyperplane $x_{d+1} = 0$ of the faces on the lower envelope of these pyramids form a refinement of the faces of the diagram $\mathcal{V}or_{L_\infty}(\mathcal{M})$. The complexity of the Voronoi diagram $\mathcal{V}or_{L_\infty}(\mathcal{M})$ is thus bounded by the complexity of a lower envelope of n simplices in $\mathbb{E}^{d+1}$:

$$|\mathcal{V}or_{L_\infty}(\mathcal{M})| = O(n^d \alpha(n)).$$

This bound is almost tight for certain sets of points that are not in general position (see exercise 18.9). If the points are in so-called L_∞-general position, then it is possible to show that the complexity of Voronoi diagrams for the L_∞ metric is the same as that for Euclidean Voronoi diagrams, namely $O(n^{\lceil d/2 \rceil})$ (see exercise 18.10). We show this for the case $d = 2$. When $d = 2$, $\mathcal{V}or_{L_\infty}(\mathcal{M})$ can

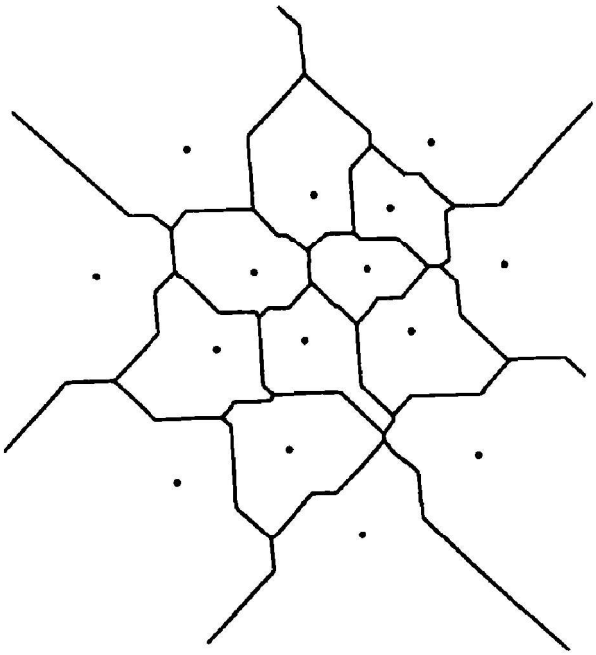

Figure 18.9. Diagram for the L_∞ metric.

be identified with the diagram $\mathcal{V}or_{L_1}(\mathcal{M})$ studied previously by simply rotating the coordinate system by an angle of $\frac{\pi}{4}$. The points are in *L_∞-general position*, if no two points are connected by a line parallel to the coordinate axes, and no four points belong to a common cube whose facets are parallel to the coordinate axes. If so, then the complexity of $\mathcal{V}or_{L_\infty}(\mathcal{M})$ is $O(n)$ in dimension 2, and the algorithm described in subsection 16.3.3 that computes the lower envelope of triangles can be used to compute this diagram in time $O(n \log^2 n)$ (see exercise 19.2 for a better algorithm).

The distances considered here, L_1 and L_∞, are particular cases of polyhedral distances, so-called because their unit ball is a polytope. Voronoi diagrams for polyhedral distances are studied in exercise 19.3.

18.5 Voronoi diagrams in hyperbolic spaces

18.5.1 Pencils of spheres

A *pencil of spheres* in $\mathbb{E}^d$ is a set $\mathcal{S}$ of spheres that are affine combinations of two given spheres Σ_1 and Σ_2:

$$\mathcal{F} = \{\Sigma : \exists \lambda \in \mathbb{R},\ \forall X \in \mathbb{E}^d,\ \Sigma(X) = \lambda \Sigma_1(X) + (1-\lambda)\Sigma_2(X)\}.$$

If we apply to spheres the mapping ϕ introduced in section 17.2, we map the spheres in $\mathbb{E}^d$ to points in $\mathbb{E}^{d+1}$. From the results of section 17.2, it follows that the image under ϕ of a pencil $\mathcal{F}$ is the line $\phi(\mathcal{F})$ in $\mathbb{E}^{d+1}$ that connects the points $\phi(\Sigma_1)$ and $\phi(\Sigma_2)$.

We may distinguish between four kinds of pencils, according to whether the line that is the image under ϕ of the pencil intersects the paraboloid in one point (transversally), in two points, is tangent to $\mathcal{P}$, or does not intersect $\mathcal{P}$ (see figure 18.10).

- If $\phi(\mathcal{F})$ intersects $\mathcal{P}$ transversally in only one point, then $\mathcal{F}$ contains a single sphere of zero radius, and $\phi(\mathcal{F})$ is a *pencil of concentric spheres.*
- If the line $\phi(\mathcal{F})$ intersects $\mathcal{P}$ in two points, $\mathcal{F}$ contains two spheres of radius zero, called the *limit points* of the pencil.
- If the line $\phi(\mathcal{F})$ is tangent to $\mathcal{P}$, then $\mathcal{F}$ may be considered as a pencil whose two limit points are identical, or as a pencil supported by a sphere that reduces to a point. Such a pencil is called a *tangent pencil.*
- If the line $\phi(\mathcal{F})$ does not intersect $\mathcal{P}$, there exists a family of hyperplanes tangent to $\mathcal{P}$ that contain $\phi(\mathcal{F})$. Let $\phi(\Sigma_\mathcal{F})$ be the set of points of $\mathcal{P}$ at which these hyperplanes are tangent to $\mathcal{P}$. Then $\phi(\Sigma_\mathcal{F})$ is the image under ϕ of the set $\Sigma_\mathcal{F}$ of points that belong to all the spheres in the pencil $\mathcal{F}$. Coming back to the definition of a pencil, we have $\Sigma(X) = 0$ for all values of λ, and this implies that $\Sigma_1(X) = \Sigma_2(X) = 0$ and that $\Sigma_\mathcal{F}$ can be identified with the $(d-1)$-sphere $\Sigma_1 \cap \Sigma_2$. All the d-spheres in the pencil $\mathcal{F}$ intersect along the $(d-1)$-sphere obtained as the intersection of any two spheres in the pencil. For this reason, $\Sigma_\mathcal{F}$ is called the *supporting sphere* of the pencil.

The very definition of a pencil of spheres implies that any point in the radical hyperplane H_{12} of two spheres Σ_1 and Σ_2 in the pencil has same power with respect to any sphere in the pencil. We may therefore define the *radical hyperplane of a pencil of spheres* as the radical hyperplane of any two spheres in the pencil.

The radical hyperplane of a pencil supported by a sphere is the affine hull of the supporting sphere. A concentric pencil has no radical hyperplane. The radical hyperplane of a pencil with limit points is the perpendicular bisector of these two points. The radical hyperplane of a tangent pencil is the hyperplane tangent to all the spheres in the pencil.

18.5.2 Voronoi diagrams in hyperbolic spaces

The Poincaré model of the hyperbolic space of dimension d is the half-space $\mathbb{H}^d = \{X \in \mathbb{E}^d \ : \ x_d > 0\}$. We will not define the hyperbolic distance precisely.

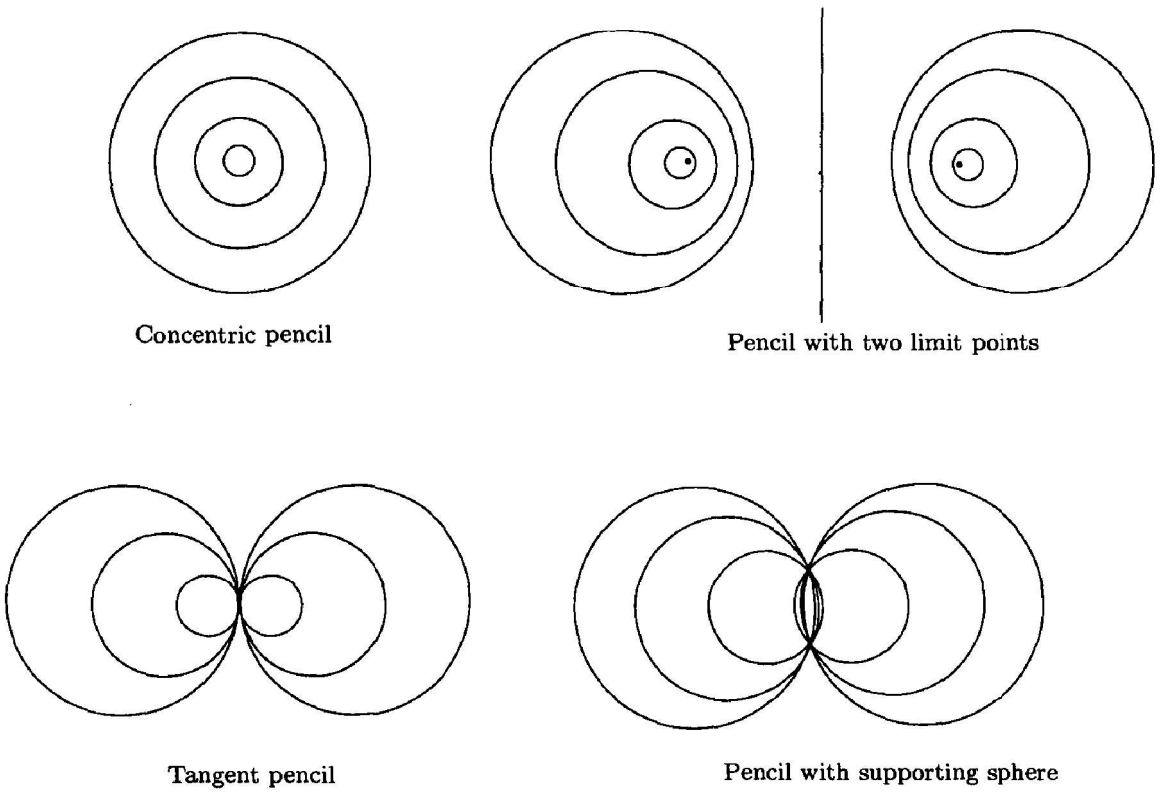

Figure 18.10. The four kinds of pencils.

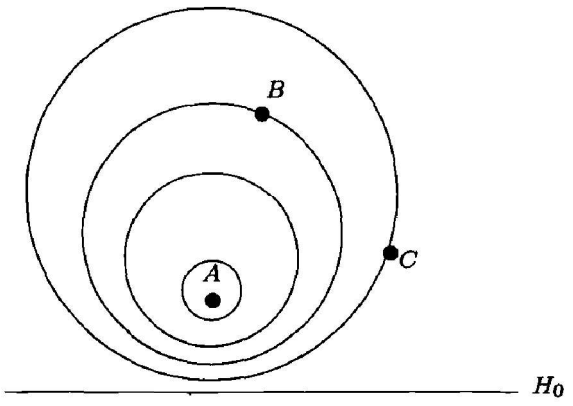

Figure 18.11. B is closer to A for the hyperbolic distance than C is.

The interested reader will find a more precise account in the classical references on the topic ([22] for instance). To define the hyperbolic diagram, it suffices to decide, given three points A, B, and C in $\mathbb{H}^d$, whether B or C is closer to A. For this, we consider the pencil $\mathcal{F}_A$ of spheres with limit points A and A', where A' denotes the symmetric of A with respect to the hyperplane H_0 of equation

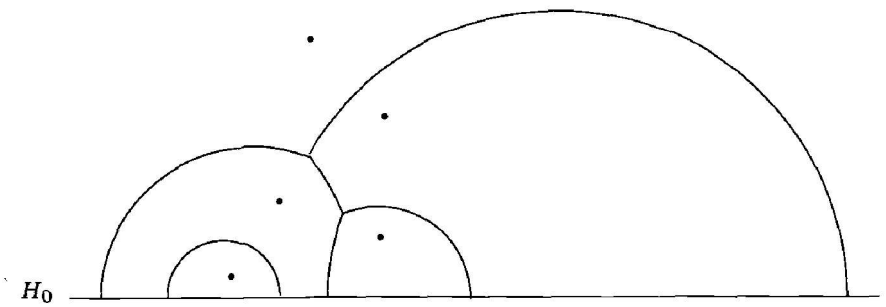

Figure 18.12. A hyperbolic Voronoi diagram in the Poincaré half-space.

$x_d = 0$. (Note that H_0 is the radical hyperplane of $\mathcal{F}_A$.) We say that B is closer to A for the hyperbolic distance if the sphere in $\mathcal{F}_A$ that passes through B has a smaller radius than the sphere in $\mathcal{F}_A$ that passes though C (see figure 18.11).[1]

Given a set $\mathcal{M} = \{M_1, \ldots, M_n\}$ of n points in the Poincaré half-space $\mathbb{H}^d$, there corresponds a region $V_h(M_i)$ in $\mathbb{H}^d$ to each point M_i in $\mathcal{M}$. This region consists of the points in $\mathbb{H}^d$ that are closer to M_i than to any other point in $\mathcal{M}$:

$$V_h(M_i) = \{X \in \mathbb{H}^d, \delta_h(X, M_i) \leq \delta_h(X, M_j) \quad \text{for any} \quad j \neq i\},$$

The *Voronoi diagram for the hyperbolic distance* of $\mathcal{M}$, also called the hyperbolic diagram of $\mathcal{M}$, is the subdivision of the Poincaré half-space induced by the equivalence relation shared by the points that have the same nearest neighbors for the hyperbolic distance. The faces of the diagram are the closures of the equivalence classes. The $V_h(M_i)$'s form the cells of the diagram (see figure 18.12).

$V_h(M_i)$ is the set of points $X \in \mathbb{H}^d$ that have M_i as a nearest neighbor. Since the locus of points in $\mathbb{H}^d$ at a given hyperbolic distance from a given point $A \in \mathbb{H}^d$ is a sphere of the pencil $\mathcal{F}_A$, it follows that, for any point X in $V_h(M_i)$, the interior of the sphere in the pencil $\mathcal{F}_X$ that passes through M_i contains no point of $\mathcal{M}$.

We can also embed $\mathbb{H}^d$ into $\mathbb{E}^{d+1}$ by identifying it with the half-hyperplane $x_{d+1} = 0$, $x_d > 0$. The hyperplane H_0 is therefore identified with the subspace $\{x_{d+1} = x_d = 0\}$. The pencil $\mathcal{F}_X$ is mapped by ϕ into a line in $\mathbb{E}^{d+1}$ parallel to the x_d-axis. Indeed, if X' is the symmetric of X with respect to H_0, the pencil $\mathcal{F}_X$ has limit points at X and X' that are mapped by ϕ to $\phi(X)$ and to $\phi(X')$, and both these images are symmetric with respect to the hyperplane $x_d = 0$ in $\mathbb{E}^{d+1}$. This implies that a point X belongs to $V_h(M_i)$ if and only if the ray parallel to the

[1] It is tempting to define the hyperbolic distance from A to a point B as the radius of the sphere in $\mathcal{F}_A$ that passes through B. This "distance" is not symmetric, however, and is not the true hyperbolic distance defined for instance in [22]. Nevertheless, in what follows, taking the pseudo-distance to be this radius or indeed any other increasing function of this radius leads to the same diagram.

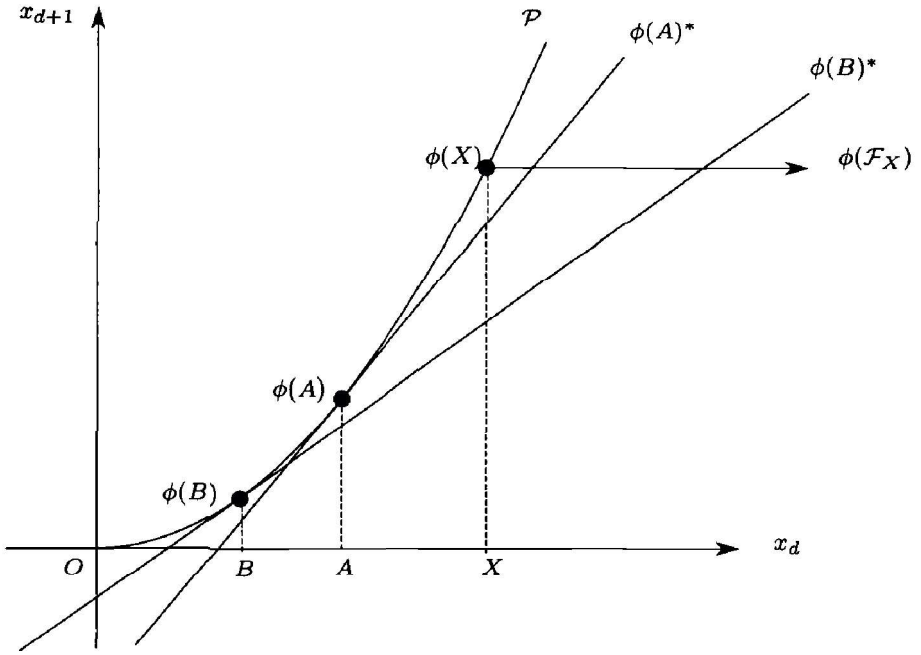

Figure 18.13. X belongs to $V_h(A)$ (d=1).

x_d-axis in $\mathbb{E}^{d+1}$ originating at $\phi(X)$ (which is entirely contained in the paraboloid $\mathcal{P}$) and directed towards $x_d > 0$ intersects the hyperplane $\phi(M_i)^*$ polar to $\phi(M_i)$ before any of the other polar hyperplanes $\phi(M_j)^*$ (see figure 18.13).

This observation has important consequences:

1. The *bisecting surface* of two points for the hyperbolic distance is a half-sphere: indeed, a point X is equidistant from A and B if and only if $\mathcal{F}_X$ contains a sphere that passes through A and B, that is, if and only if $\phi(\mathcal{F}_X)$ intersects $\phi(A)^* \cap \phi(B)^*$. In other words, $\phi(X)$ belongs to Γ, the projection of $\phi(A)^* \cap \phi(B)^*$ parallel to the x_d-axis onto the paraboloid $\mathcal{P}$ (more exactly, the half of the paraboloid that is in the half-space $x_d > 0$). But Γ is the intersection of $\mathcal{P}$ with a hyperplane H in $\mathbb{E}^{d+1}$ parallel to the x_d-axis. Its vertical projection onto $x_{d+1} = 0$ is a sphere Σ_{AB} (lemma 17.2.2), centered on H_0 by symmetry. Moreover, Σ_{AB} belongs to the pencil with limit points A and B. Indeed, the spheres in this pencil are mapped by ϕ to the points

$$\phi(\Sigma) = \lambda\phi(A) + (1 - \lambda)\phi(B).$$

The corresponding polar hyperplanes have equations

$$\phi(\Sigma)^* = \lambda\phi(A)^* + (1 - \lambda)\phi(B)^*,$$

and they are all the hyperplanes that contain $\phi(A)^* \cap \phi(B)^*$. H is thus a hyperplane polar to a sphere in the pencil $\mathcal{F}_{AB}$ that has two limit points A and B.

But H is the hyperplane polar to $\phi(\Sigma_{AB})$ (see lemma 17.2.2). As a result, Σ_{AB} belongs to $\mathcal{F}_{AB}$. Finally, Σ_{AB} is the unique sphere in $\mathcal{F}_{AB}$ that is centered on H_0.

2. A point X is equidistant from $d+1$ points $A_0, \ldots, A_d$ if and only if $\phi(X)$ is the projection of $\bigcap_{i=0}^{d} \phi(A_i)^*$ parallel to the x_d-axis, onto the half-paraboloid. The point at equal hyperbolic distance from $d+1$ points is the limit point of the pencil that contains the sphere circumscribed to the $d+1$ points of radical hyperplane H_0.

3. The hyperbolic Voronoi diagram can be obtained by projecting the polytope $\mathcal{V}(\mathcal{M}) = \bigcap_{i=1}^{n} \phi(A_i)^{*+}$ parallel to the x_d-axis onto the half-paraboloid, then projecting the result vertically onto the hyperplane $x_{d+1} = 0$. Note that the projection parallel to the x_d-axis does not map all the points of $\mathcal{V}(\mathcal{M})$ onto the half-paraboloid. This double projection establishes an injective correspondence between the Euclidean and the hyperbolic Voronoi diagrams of $\mathcal{M}$. More directly, these two projections can be avoided by performing the single following transformation. Replace the planar $(d-1)$-faces of the Euclidean diagram that are (at least partly) contained in the half-space $x_d > 0$, by the corresponding portions of spheres (hyperbolic bisectors limited to $x_d > 0$); a k-face ($k < d-1$) of the Euclidean diagram is the intersection of $d-k+1$ planar $(d-1)$-faces, and is replaced by the portion of surface that is the intersection of the $d-k+1$ corresponding spherical faces. From the injective correspondence between Euclidean and hyperbolic diagrams, we deduce the following theorem.

Theorem 18.5.1 *The complexity of the hyperbolic Voronoi diagram of n points in the hyperbolic Poincaré half-space $\mathbb{H}^d$ is $\Theta(n^{\lceil d/2 \rceil})$. Such a diagram can be computed in time $\Theta(n \log n + n^{\lceil d/2 \rceil})$.*

18.6 Exercises

Exercise 18.1 (Greatest empty rectangle) Let X and A be two points in $\mathbb{E}^2$. The quadratic distance $\delta_Q(X, A)$ is defined as

$$\delta_Q(X, A) = (X - A)\Delta(X - A) \quad \text{with} \quad \Delta = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Show that $\delta_Q(X, A)$ is the area of the rectangle whose sides are parallel to the coordinate axes and of which A and X are two opposite vertices. Given a set $\mathcal{S}$ of points in the plane, show that its diagram for this quadratic distance function can be used to compute the rectangle of greatest area whose sides are parallel to the coordinate axes, whose sides each contain at least one point of $\mathcal{S}$, and whose interior does not contain any point of $\mathcal{S}$.

Hint: To find a greatest empty rectangle, use a divide-and-conquer algorithm. The merge step consists in finding the greatest empty rectangle intersected by the separating

line. The greatest empty rectangle for which three points of contact lie on one side of the separating line (and the fourth on the other side) can be found easily. Two points of contact on one side of the separating line define a corner of the rectangle, and the corners of empty rectangles are the so called maxima and can be found in $O(n \log n)$ time. The greatest empty rectangle with two points of contact on either side of the separating line is defined by an opposite pair (A, B) of maxima such that the segment connecting them is an edge of the affine diagram defined for the generalized quadratic distance $\delta_Q(A, B)$. The complexity of the merge step is $O(n \log n)$, hence the total algorithm runs in time $O(n \log^2 n)$.

Exercise 18.2 (Lower envelope of cones) Show that the lower envelope of n vertical cones of revolution in $\mathbb{E}^d$ has complexity $O(n^{\lfloor d/2 \rfloor + 1})$ and can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$. If the vertices of the cones are all contained in a given horizontal hyperplane, and if their angles are all identical, then the complexity of the lower envelope drops to $O(n^{\lfloor \frac{d+1}{2} \rfloor})$ and it can be computed in time $O(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor})$.

Exercise 18.3 (Spheres and disks) According to the general definition of Voronoi diagrams, we may define the Voronoi diagram of a set of disks $D_1, \ldots, D_n$ as usual, where the distance of a point X from a disk D_i centered at C_i and of radius r_i is defined by

$$d_S(X, D_i) = \max(0, \|XC_i\| - r_i).$$

Show that the Voronoi diagram of n disks in $\mathbb{E}^d$, where $d \geq 3$, has complexity $O(n^{\lfloor d/2 \rfloor + 1})$ and that it can be computed in time $O(n^{\lfloor d/2 \rfloor + 1})$. If $d = 2$, show that these bounds are $O(n)$ and $O(n \log n)$ respectively.

Hint: To each C_i, give a weight r_i and compute the diagram of the disks knowing the additive diagram of their centers. In the discussion of subsection 18.3.1, the cone $\mathcal{C}_i$, $i = 1, \ldots, n$ must be replaced by the same cone truncated by the halfspace $x_{d+1} \geq 0$.

Exercise 18.4 (Diagrams with multiplicative weights) Show that $\Omega(n^{\lfloor d/2 \rfloor + 1})$ is a lower bound on the complexity of the Voronoi diagram of n points in $\mathbb{E}^d$ with multiplicative weights.

Hint: Generalize the example of figure 18.5.

Exercise 18.5 (Regular complex) Let $\mathcal{C}$ be a d-complex in $\mathbb{E}^d$. We say that $\mathcal{C}$ is *regular* if it can be obtained as the vertical projection of a polytope in $\mathbb{E}^{d+1}$. Show that $\mathcal{C}$ is regular is and only if it is a power diagram. Show that any simple complex (meaning that its cells all consist of simple polytopes) is regular (the best-known examples are arrangements of hyperplanes in general position). Devise an algorithm that determines whether a complex is regular and, if so, computes the corresponding polyhedron.

Hint: Use theorem 18.2.1. For hyperplane arrangements, the connection with zonotopes is particularly helpful (see exercise 14.8).

Exercise 18.6 (The inverse problem) Show that it is possible to determine whether a complex is a Voronoi diagram and, if so, to compute the corresponding sites in time linear in the total complexity of its cells.

Exercise 18.7 (Spider webs) By a *spider web*, we mean the 1-skeleton of a 2-complex that covers $\mathbb{E}^2$. Show that if the spider web is the skeleton of a power diagram, then we can assign a tension to each edge such that each vertex is in an equilibrium state.

Hint: For the tension of an edge, take the length of the dual edge. An edge and its dual edge are perpendicular, and the dual edges of the edges incident to a vertex S form a cycle that we orient counter-clockwise. At a vertex S, the sum of the tensions equals the sum of the vectors of the dual edges, so that the total tension vanishes at the vertices.

Exercise 18.8 (Cubes and co-cubes) Show that in $\mathbb{E}^3$, several homothetic cubes or co-cubes may pass through four points even though these points are in L_∞-general position.

Exercise 18.9 (Degenerate positions for L_1 and L_∞ distances) Show that in $\mathbb{E}^2$, the Voronoi diagram for the L_1 metric of points that are along one of the main bisectors is quadratic. Show that if the bisector of two points on a line parallel to one of the main bisectors is redefined as the Euclidean perpendicular bisector, then the complexity of the diagram becomes linear, and a cell is formed by the set of points that share exactly one common nearest neighbor for the L_1 distance (but do not necessarily have the same subset of nearest neighbors). Generalize the example above to show that $\Omega(n^d)$ is a lower bound on the complexity of a Voronoi diagram of n points in $\mathbb{E}^d$ for the L_1 metric. Also give similar results for the L_∞ metric.

Exercise 18.10 (Complexity of $\mathcal{V}or_{L_\infty}$) Show that the complexity of a Voronoi diagram for the L_∞ metric of a set $\mathcal{M}$ of n points in $\mathbb{E}^d$ in L_∞-general position is $O(n^{\lceil d/2 \rceil})$.

Hint: It suffices to bound the number of so-called maximal placements of a maximal cube whose facets are perpendicular to the coordinate axes, and whose interior contains no point of $\mathcal{M}$. A contact is a pair formed by a facet of such a cube and by a point in $\mathcal{M}$. A placement realizes a contact of multiplicity k at a point if this point belongs to k facets of the corresponding cube. If $k = 1$, the contact is said to be simple. For a given maximal placement, the sum of the multiplicities of the points of contact is $d + 1$. First show that any maximal placement realizes two contacts, called parallel contacts, whose facets are parallel. We say that a maximal placement is reducible if at least one of its parallel contacts is simple and if the other does not have multiplicity d. Show then that it suffices to bound the number of irreducible maximal placements. For this, charge a reducible placement to an irreducible one by applying the following procedure as many times as needed: Scale up the cube by a homothety centered at one of its vertices that lies on a facet involved in some parallel contact. In this way we obtain a smaller cube contained in the preceding one but whose multiplicity is increased for at least one of the contacts. Show that an irreducible placement is charged by at most $O(1)$ reducible placements. Finally show that the number of irreducible placements is $O(n^{\lceil d/2 \rceil})$. For this, notice that the centers of such placements belong to some affine subspace of dimension at most $d - 3$. In this subspace, the centers of maximal placements correspond to the vertices of a union of n cubes of same size, so we may use the result of exercise 4.8.

Exercise 18.11 (Simplicial distance) Let S be a $(d+1)$-simplex in $\mathbb{E}^d$ that contains the origin O. We denote by λS the image of this polytope under the homothety centered at O and of ratio λ. The simplicial distance $\delta_S(X, A)$ from point X to point A is defined as the smallest real $\lambda \geq 0$ such that $X - A$ belongs to λS. Show that the complexity of a Voronoi diagram for a simplicial distance of a set $\mathcal{M}$ of n points in $\mathbb{E}^d$ is $O(n^{\lceil d/2 \rceil})$.

Hint: We define a reducible placement of $\mathcal{S}$ as in exercise 18.10: it is a placement that has several simple contacts. The number of irreducible placements is $O(n^{\lceil d/2 \rceil})$ and we can also show that the same bound holds for the reducible placements.

Exercise 18.12 (Hyperbolic bisector) Show that the equation of the hyperbolic bisector Σ_{AB} of two points $A = (a_1, \ldots, a_d)$ and $B = (b_1, \ldots, b_d)$ is

$$(a_d - b_d)X \cdot X + 2(b_d A - a_d B) \cdot X - b_d A \cdot A + a_d B \cdot B = 0.$$

Exercise 18.13 (The Poincaré disk) Rather than using the Poincaré half-space $\mathbb{H}^2$ as a model of the hyperbolic space, we introduce the Poincaré disk $\mathbb{D}$ which can be derived from $\mathbb{H}^2$ by a homographic transformation. More precisely, if the Poincaré half-space is identified with the complex half-plane $\{z \in \mathbb{C}, \mathrm{Im} z > 0\}$, the homographic map defined by

$$h(z) = \frac{z - i}{z + i}$$

is a bijection from $\mathbb{H}^2$ into $\mathbb{D}$. Show that the edge that joins two points remains a circle centered on the boundary of $\mathbb{D}$, and also that the points at equal distance from A are on a circle that belongs to the pencil that has A as a limit point and that contains the boundary of $\mathbb{D}$. From this, explain how to compute the Voronoi diagram of a set of points in $\mathbb{H}^2$.

Exercise 18.14 (Dual of a hyperbolic diagram) Show that we may dualize the hyperbolic Voronoi diagram of a set of points $\mathcal{M}$ in $\mathbb{H}^d$ by projecting the convex hull of $\phi(\mathcal{M})$ parallel to the x_d-axis onto the half-paraboloid, and then projecting the result of this first projection onto the hyperplane $x_{d+1} = 0$. Show that this dual is in bijection with a sub-complex of the Delaunay complex.

18.7 Bibliographical notes

Power diagrams were studied by Aurenhammer [14] and by Imai, Iri, and Murota [129]. Affine diagrams are defined in [17] by Aurenhammer and Imai, who also show their connection with power diagrams and diagrams with additive and multiplicative weights. Solutions to exercises 18.5 and 18.6 are due to Aurenhammer. The solution to exercise 18.1 is adapted from that given by Chazelle, Drysdale, and Lee [47] and Aurenhammer. Spider webs were already analyzed by Maxwell in the nineteenth century. Recent references can be found in the article by Ash, Bolker, Crapo, and Whiteley [12].

Diagrams for the L_1 and L_∞ metrics in the plane were studied by Lee and Wong [149] then by Lee and Drysdale [147]. The generalization to general convex distances is tackled

by Chew and Drysdale [61]. Voronoi diagrams for the L_1 and L_∞ metrics in dimensions 3 and higher (see exercise 18.10) and also simplicial distances (see exercise 18.11) are treated by Boissonnat, Sharir, Tagansky, and Yvinec [34]. In the plane, Klein proposes a notion of abstract Voronoi diagram [139] and Klein, Mehlhorn, and Meiser describe a randomized algorithm that computes such diagrams [141].

Diagrams for the hyperbolic distance are studied by Boissonnat, Cérézo, Devillers, and Teillaud [26], who present an application to shape reconstruction from plane sections.