Chapter 3
Weighted Diagrams

Every region in a 2-dimensional Voronoi diagram consists of all points for which the corresponding site minimizes the Euclidean distance. In this section, we modify the notion of distance by introducing weights. Particular attention will be paid to the case of subtracting the weight from the squared Euclidean distance because this gives convex regions, like for the unweighted Euclidean distance.

3.1 Apollonius Diagrams

Let S be a finite set of points or *sites* in the plane. For each $s \in S$, we let $w_s \in \mathbb{R}$ be its weight. Suppose we define the *weighted distance* of a point $x \in \mathbb{R}^2$ from s by subtracting the weight from the Euclidean distance:

$$d_A(s, x) \;=\; \|x - s\| - w_s. \tag{3.1}$$

Drawing the weighted site as a circle with center s and radius $|w_s|$, we can interpret this notion of distance geometrically. If $w_s \geq 0$ and $\|x - s\| \geq w_s$, then $d_A(x, s)$ is the Euclidean distance to the nearest point on the circle. The same interpretation works for $\|x - s\| < w_s$ except that $d_A(s, x)$ is now negative. Finally, if $w_s < 0$, then $d_A(s, x)$ is the Euclidean distance to the furthest point on the circle. To unify the three cases, we draw a vertical cone in $\mathbb{R}^3$, adding a third coordinate to the plane. Its axis of rotation is the vertical line passing through s, its apex is the point $(s, -w_s) \in \mathbb{R}^3$, and its opening angle is 90°; see Fig. 3.1. The weighted distance from x to s is then the vertical distance from x to the cone. In other words, (the surface bounding) the cone is the graph of the function that maps x to the weighted distance from s. Let now s and t be two weighted sites. The *bisector* consists of all points x with equal weighted distance from both: $d_A(s, x) = d_A(t, x)$ or, equivalently, $\|x - s\| - \|x - t\| = w_s - w_t$. This is the equation of a hyperbola; it is the vertical projection of the intersection of the two cones to $\mathbb{R}^2$. On one side of the bisector,

H. Edelsbrunner, *A Short Course in Computational Geometry and Topology*,
SpringerBriefs in Mathematical Methods, DOI: 10.1007/978-3-319-05957-0_3,

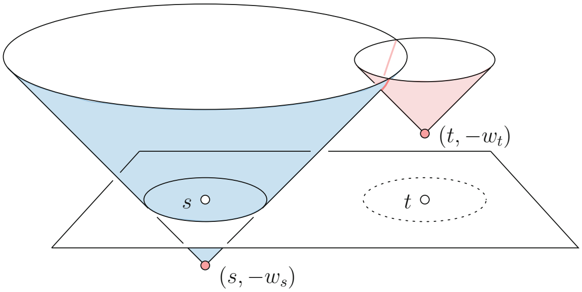

Fig. 3.1 The cones of two sites in the plane, one with positive and the other with negative weight

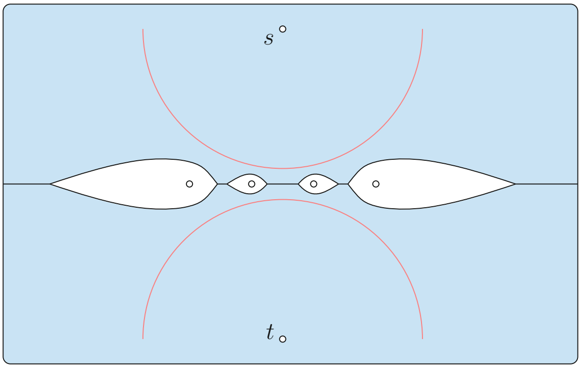

Fig. 3.2 The Apollonius diagram of six sites. The intersection of the regions of s and t consists of five segments

we have $d_A(s, x) < d_A(t, x)$, and on the other side we have the reverse inequality. Returning to the set of finitely many sites with real weights, the region in which s minimizes the weighted distance is

$$A_s = \{x \in \mathbb{R}^2 \mid d_A(s, x) \leq d_A(t, x), \quad \forall t \in S\}, \tag{3.2}$$

and the *Apollonius diagram* of S is the set of such regions. In contrast to the Voronoi diagram, the regions in the Apollonius diagram are not necessarily convex; only the region of the site with the smallest weight is guaranteed to be convex. Nevertheless, every non-empty region is connected. In contrast, the intersection of two regions is not necessarily connected; see Fig. 3.2. Recall how we defined the Delaunay triangulation by connecting sites whose Voronoi regions have a non-empty intersection. Because of the more complicated intersections, this construction is no longer as straightforward for the Apollonius diagram.

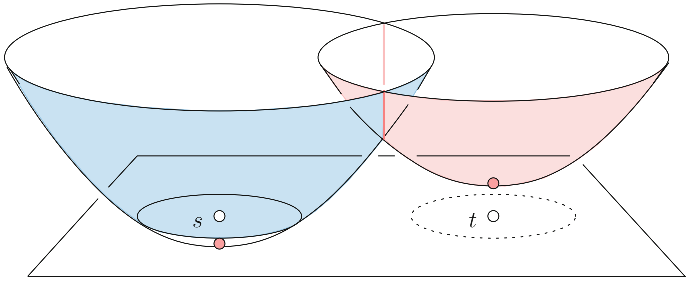

Fig. 3.3 The paraboloids of two sites in the plane, one with positive and the other with negative weight

3.2 Power Diagrams

Instead of subtracting the weight from the Euclidean distance, we now subtract it from the square of the Euclidean distance:

$$d_P(s, x) \;=\; \|x - s\|^2 - w_s. \tag{3.3}$$

Here, $d_P(s, x)$ is called the *power* or *power distance* of x from s. Recall that the bisector consists of all points that satisfy $d_P(s, x) = d_P(t, x)$. Writing the squared Euclidean distance as a scalar product, we get

$$\|x - s\|^2 \;=\; \langle x - s, x - s\rangle \;=\; \|x\|^2 - 2\langle x, s\rangle + \|s\|^2.$$

The equation of the bisector is therefore $2\langle x, t - s\rangle = \|t\|^2 - \|s\|^2 + w_s - w_t$, which is the equation of a line. We can see this geometrically, by drawing the graph of the function that maps a point to its power distance from s; it is a paraboloid with vertical axis of rotation and lowest point at $(s, -w_s)$. We also draw the graph of the function for t, which is a translate of the paraboloid for s; see Fig. 3.3. The intersection of the two paraboloids is a parabola that lies in a vertical plane; its projection to $\mathbb{R}^2$ is a line. Returning to the set of finitely many sites with weights, the region within which s minimizes the power distance is

$$P_s \;=\; \{x \in \mathbb{R}^2 \mid d_P(s, x) \le d_P(t, x), \quad \forall t \in S\}, \tag{3.4}$$

and the *power diagram* of S is the set of such regions. Since the bisectors are straight lines, each region is the intersection of half-planes and therefore convex. In contrast to the unweighted case, a site may have an empty region in the power diagram. Otherwise, the power diagrams are visually difficult to distinguish from Voronoi diagrams, which is perhaps the reason why they are sometimes referred to as *weighted Voronoi diagrams*.

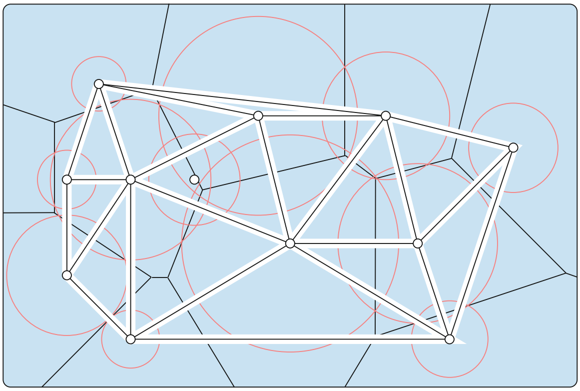

Fig. 3.4 The power diagram and the superimposed weighted Delaunay triangulation of twelve sites in the plane. Note that only eleven of the sites function as vertices of the triangulation

3.3 Weighted Delaunay Triangulations

Having convex regions, we can again draw the dual, connecting two sites with an edge whenever the corresponding two power regions share an edge; see Fig. 3.4. Similar to the unweighed case, this construction typically gives a triangulation, but unlike the unweighted case, not every site is necessarily also a vertex in that triangulation. To understand the construction, it will be useful to generalize the circumscribed circles that characterize the triangles in the unweighted Delaunay triangulation. To this end, we say a point x with weight w_x is *orthogonal* to s if

$$\|x - s\|^2 \;=\; w_s + w_x. \tag{3.5}$$

Drawing s as the circle with center s and radius $\sqrt{w_s}$, and similarly for x, we get two circles that meet at a right angle.[1] Given s, w_s, and x, we can always find w_x such that the two weighted points are orthogonal. In fact, the weight for which this is the case is unique. However, if in addition we are given t and w_t, we can find w_x such that x is orthogonal to s as well as to t only if x lines on the bisector of s and t. Finally, if in addition we are given u and w_u, then we can find w_x such that x is orthogonal to all three only if x lies on all three bisectors or, equivalently, it is the unique point at which the three regions in the power diagram of s, t, u meet.

We can now formulate a criterion for stu being a triangle in the weighted Delaunay triangulation of S. In the unweighted case, we said S is in general position if no four sites lie on a common circle. The appropriate notion in the weighted case is that no four unweighted sites have a common orthogonal circle.

[1] For this to be true, we have to assume that both weights are positive. A similar but geometrically not quite as compelling interpretation can also be found for the case in which one of the weights is non-positive.

Lemma A *Let S be a finite set of weighted sites in general position in $\mathbb{R}^2$. Then $s, t, u \in S$ form a triangle in the weighted Delaunay triangulation of S iff $\|x - v\|^2 \geq w_v + w_x$ for every $v \in S$, where x is the unique weighted point orthogonal to s, t, and u.*

3.4 Geometric Primitives

There are many algorithms for weighted and unweighted Delaunay triangulations, one being the incremental construction briefly mentioned at the end of last lecture. All these algorithms need to be able to decide whether the circle that passes through three sites is empty (see Lemma B (i) in Chap. 2) or, in the weighted case, whether the circle orthogonal to three weighted sites is further than orthogonal from all other sites (see Lemma A). We now study the details of these decisions, beginning with the unweighted case.

Let $a = (a_1, a_2)$, $b = (b_1, b_2)$, $c = (c_1, c_2)$, be three points in $\mathbb{R}^2$. If the points lie on a common line, then we can write $c = (1-\lambda)a + \lambda b$, assuming $a \neq b$. Hence, the determinant of

$$\Delta = \begin{bmatrix} 1 & a_1 & a_2 \\ 1 & b_1 & b_2 \\ 1 & c_1 & c_2 \end{bmatrix} \tag{3.6}$$

vanishes. Indeed, $\det \Delta = 0$ iff a, b, c are collinear. In addition, $\det \Delta > 0$ iff a, b, c form a left-turn. We can formulate a similar test for cocircularity. Lift a to the point $a^+ = (a_1, a_2, a_1^2 + a_2^2)$ in $\mathbb{R}^3$, and similarly for b and c. Furthermore, let $v = (v_1, v_2)$ be a fourth point, write $v^+ = (v_1, v_2, v_1^2 + v_2^2)$, and define

$$\Gamma = \begin{bmatrix} 1 & a_1 & a_2 & a_1^2 + a_2^2 \\ 1 & b_1 & b_2 & b_1^2 + b_2^2 \\ 1 & c_1 & c_2 & c_1^2 + c_2^2 \\ 1 & v_1 & v_2 & v_1^2 + v_2^2 \end{bmatrix}. \tag{3.7}$$

The crucial insight is that $\det \Gamma = 0$ iff v lies on the circle determined by a, b, c. To see this, we interpret the determinant as an expression of the points a^+, b^+, c^+, v^+ in $\mathbb{R}^3$. It vanishes iff the four points lie on a common plane. Now, we just need to verify that being coplanar in $\mathbb{R}^3$ corresponds to being cocircular in $\mathbb{R}^2$. To see this, we intersect the paraboloid given by $x_3 = x_1^2 + x_2^2$ with the plane given by $x_3 = 2Ax_1 + 2Bx_2 + C$. Eliminating x_3, we get $(x_1 - A)^2 + (x_2 - B)^2 - (A^2 + B^2 - C) = 0$, which is the equation of a circle; see Fig. 3.5.

Lemma B *Let a, b, c, v be points in $\mathbb{R}^2$ and Δ, Γ the matrices defined in* (3.6) *and* (3.7). *Then the point v belongs to the open disk bounded by the circle passing through a, b, c iff* $\det \Delta \cdot \det \Gamma < 0$.

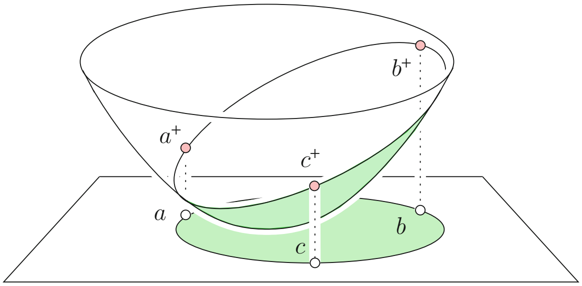

Fig. 3.5 Lifting a, b, c to the paraboloid, we get the circle that passes through them by projecting the intersection between the paraboloid and the plane passing through the lifted points

We omit the proof, which consists of two parts: verifying that v lies on the circle if $\det \Gamma = 0$, and making sure that the sign is negative inside and positive outside the circle. We note that the extra factor, $\det \Delta$, is necessary because switching two points of a, b, c changes the sign of the determinant without changing the geometric configuration. The test explained in Lemma B can be generalized to the weighted case by changing the matrix in (3.7) to

$$\Gamma_W = \begin{bmatrix} 1 & a_1 & a_2 & a_1^2 + a_2^2 - w_a \\ 1 & b_1 & b_2 & b_1^2 + b_2^2 - w_b \\ 1 & c_1 & c_2 & c_1^2 + c_2^2 - w_c \\ 1 & v_1 & v_2 & v_1^2 + v_2^2 - w_v \end{bmatrix}. \tag{3.8}$$

Indeed, if we project the intersection of the paraboloid with the plane that passes through the thus lifted points of a, b, c, then we get the unique orthogonal circle. We can now modify Lemma B.

Lemma C *Let a, b, c, v be weighted points in $\mathbb{R}^2$ and Δ, Γ_W the matrices defined in (3.6) and (3.8). Letting x with w_x be the weighted point orthogonal to a, b, c, we have $\|v - x\|^2 < w_v + w_x$ iff* $\det \Delta \cdot \det \Gamma_W < 0$.

Note that Lemma C agrees with Lemma B when the weights of a, b, c, v vanish.