Chapter 5
Alpha Complexes

The original motivation for the concept of alpha shapes was the desire to develop a concrete version of the intuitive notion of 'shape' of a finite point set. Starting from this idea, we explore connections to Voronoi diagrams and Delaunay triangulations.

5.1 Jarvis' Construction

In an early approach, Jarvis defines the shape of a point set procedurally, as the output of a generalized convex hull algorithm [1]. Letting S be a finite set of sites in $\mathbb{R}^2$, we can construct the convex hull by rotating a line about the set; see Fig. 5.1. Assuming general position, we begin by letting s be the leftmost site. Drawing a vertical line L through s and oriented downward, we see that all sites lie to the left of L, so s is indeed a vertex of the convex hull. Using s as a pivot, we rotate L in a ccw order until it hits another site, t. All sites other than s and t lie to the left of L, implying that st is an edge of the convex hull. We repeat this step now using t as the pivot. Each step gives a new edge of the convex hull, and the algorithm halts when it returns to the initial site, s.

What else we can do by wrapping the set? For example, we can construct non-convex shapes if we shorten the rotating line to a line segment of fixed length [2]. This works if the sites are nicely distributed, as in Fig. 5.1, but can get the line segment lost and return a sequence of sites that cannot be reasonably called its shape. Another shortcoming is that this algorithm necessarily constructs only one closed curve, while two closed curves might be more appropriate for the example at hand; see Fig. 5.2.

5.2 The Alpha Shape

A theoretically and practically more satisfying solution to constructing the shape of a finite set can be based on empty disks [3]. Letting $\alpha \geq 0$ be a fixed radius, we write $D_x(\alpha)$ for the closed disk with center $x \in \mathbb{R}^2$ and radius α. It is *empty* if it contains

H. Edelsbrunner, *A Short Course in Computational Geometry and Topology*, SpringerBriefs in Mathematical Methods, DOI: 10.1007/978-3-319-05957-0_5,

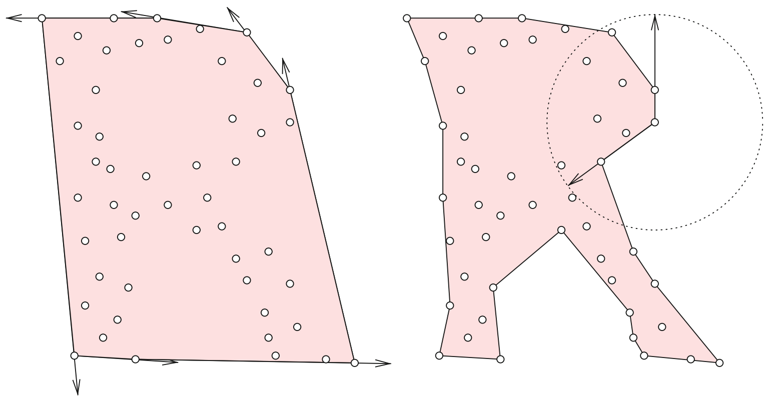

Fig. 5.1 *Left* illustration of Jarvis' convex hull algorithm. *Right* the generalization to constructing shapes obtained by shortening the line to a line segment

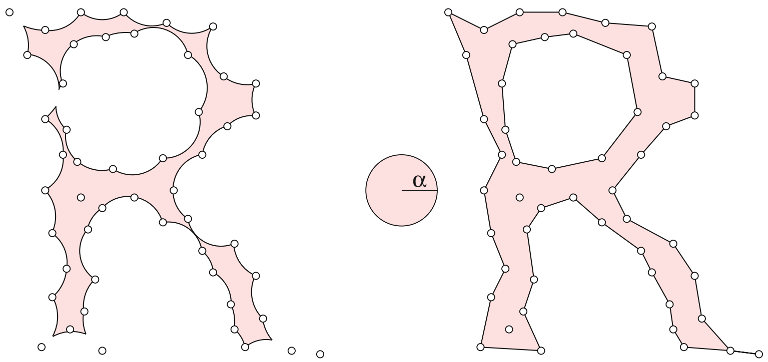

Fig. 5.2 A set of points sampling the letter 'R', with its α-hull on the *left* and its α-shape on the *right*

no site: $D_x(\alpha) \cap S = \emptyset$. The α*-hull* of S is the complement of the union of empty disks of radius α; see Fig. 5.2. Setting α to zero, we get the set of points, and setting it to infinity, we get the convex hull.

The curved edges of the α-hull can sometimes be annoying. This motivates us to draw them straight, which results in the α*-shape* of S. We will give an alternative definition shortly, which will eliminate any remaining ambiguities. In contrast to the α-hull, the α-shape is a polyhedron in the general sense: it does not have to be convex, and it can have different intrinsic dimension at different places. For example, the α-shape in Fig. 5.2 is mostly 2-dimensional except it has a 1-dimensional extension at the right leg. If S contains a site that is further than distance 2α from any other site, then this point is isolated and forms a locally 0-dimensional portion of the α-shape.

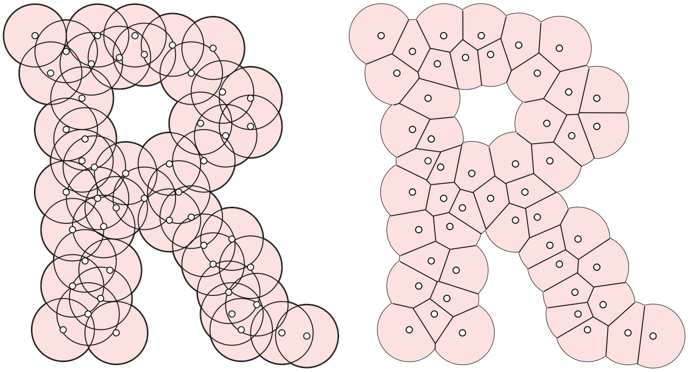

Fig. 5.3 *Left* the union of disks of radius α of the same points as in Fig. 5.2. *Right* the Voronoi decomposition of the union

5.3 Union of Disks

A point x is the center of an empty disk of radius α iff it is further than α from every site. To make this relationship concrete, we construct the union of disks of radius α centered at the sites:

$$\mathbb{U}_S(\alpha) = \bigcup_{s \in S} D_s(\alpha); \tag{5.1}$$

see Fig. 5.3. This union is the complement of the set of centers of the empty disks. We can therefore expect that the boundaries of the α-hull and the union of disks are related. Indeed, for each circular arc of $\mathbb{U}_S(\alpha)$, we have a vertex of the α-hull, and for each vertex of $\mathbb{U}_S(\alpha)$, we have a circular arc of the α-hull. In its details, this relationship is troubled by arcs in the boundary of the α-hull that intersect and partially or completely erase one another. An example are the two circular arcs that connect the right leg of the 'R' to its last site. Since they lie inside each other's empty disks, they do not belong to the α-hull.

5.4 Voronoi Decomposition

To get a cleaner relationship between the union of disks and the α-shape, we need an unambiguous definition of the latter. For this, we overlay the union of disks with the Voronoi diagram, effectively decomposing the union into convex regions; see Fig. 5.3. To formalize this idea, we write $R_s(\alpha) = V_s \cap D_s(\alpha)$ and note that this is a convex set because it is the intersection of convex sets. Furthermore, $\mathbb{U}_S(\alpha) = \bigcup_{s \in S} R_s(\alpha)$. In words, the regions R_s cover the union, but in contrast to the disks, which also cover the union, they do this without overlap. More specifically, the common intersection of the regions is limited to shared edges and vertices. Following the recipe for the

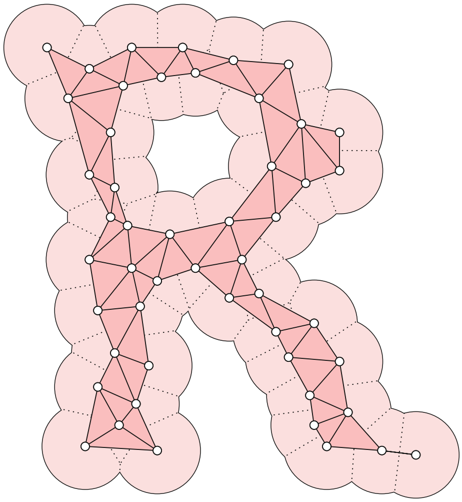

Fig. 5.4 The α-complex is superimposed on the union of disks, which is decomposed into convex regions by the Voronoi diagram

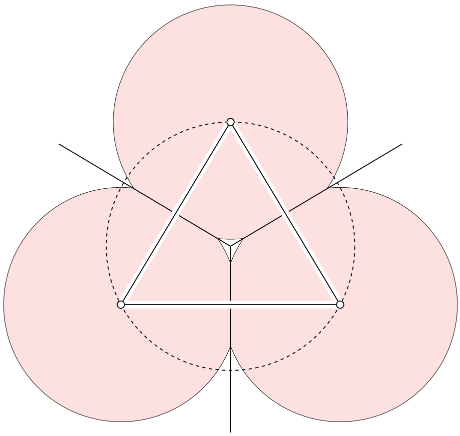

Fig. 5.5 While the intersection of any two regions is non-empty, the common intersection of all three regions is empty because α is smaller than the radius of the circumcircle. Accordingly, the α-complex contains the three edges but not the triangle spanned by the three sites

Delaunay triangulation, we construct the α*-complex* by drawing an edge between two sites if their regions intersect in a common edge, and by drawing a triangle between three sites if their regions intersect in a common point; see Fig. 5.4. For ease of reference, we denote the resulting complex by $A_S(\alpha)$, or by $A(\alpha)$ if the set of sites is understood. For example, $A(0)$ is the set of sites without any additional structure, and $A(\infty)$ is the Delaunay triangulation of S. We now formally define the α*-shape* as the union of simplices in the α-complex.

Note that the global connectivity of the union of disks in Fig. 5.4 is that same as that of the α-complex and of the α-shape: all three are connected and have a single hole. This is not a coincidence but rather a consequence of the Nerve Theorem, which will be discussed later in this course.

5.5 Filtration

Next we vary α and consider the complete range of possible values, which is from 0 to ∞. For $\alpha < \alpha'$, we have $D_s(\alpha) \subseteq D_s(\alpha')$ and therefore $R_s(\alpha) \subseteq R_s(\alpha')$. Recall that st is an edge in the α-complex iff $R_s(\alpha) \cap R_t(\alpha) \neq \emptyset$. Since the regions grow with the radius, this implies $R_s(\alpha') \cap R_t(\alpha') \neq \emptyset$, and therefore st is also an edge in the α'-complex. Similarly, every triangle in the α-complex belongs to the α'-complex. In summary, $A(\alpha) \subseteq A(\alpha')$ whenever $\alpha \leq \alpha'$. It thus makes sense to ask—for each vertex, edge, and triangle σ in the Delaunay triangulation—what the smallest value of α is for which σ belongs to $A(\alpha)$. Denoting this value by α_σ, we can construct the α-complex simply by collecting all vertices, edges, and triangles that have a value not larger than α:

$$A(\alpha) = \{\sigma \in K \mid \alpha_\sigma \leq \alpha\}, \tag{5.2}$$

where K is the Delaunay triangulation of S. Computing this value is easiest for the vertices, since we have $\alpha_s = 0$ for every $s \in S$. It is also easy for triangles, for which α_{stu} is the radius of the circumcircle; see Fig. 5.5. The computation of the smallest α-value is slightly more difficult for edges. Here, we distinguish between two cases. First, the edge st may intersect the dual Voronoi edge in its interior; as in Fig. 5.5. In this case, st belongs to the α-complex as soon as the two disks meet, which happens when the radius reaches half the distance between the sites: $\alpha_{st} = \frac{1}{2}\|s - t\|$. If st is shared by the triangles stu and stv in K, then the condition of st intersecting the dual Voronoi edge in its interior is equivalent to having acute angles at u and v. This leads us to the second case in which one of these two angles is obtuse, say the angle at u. Then it is not enough that the two disks meet; they also need to reach the Voronoi edge, which happens when the triangle stu enters the α-complex. Hence, $\alpha_{st} = \alpha_{stu}$. Now that we have the threshold value for every vertex, edge, and triangle in the Delaunay triangulation, we can sort them such that

$$\alpha_{\sigma_1} \leq \alpha_{\sigma_2} \leq \ldots \leq \alpha_{\sigma_n}. \tag{5.3}$$

The corresponding sequence of simplices is called a *filter*. Here, we make sure that every simplex is preceded by its faces. We get this from the formulas already, since the value of every edge is smaller than or equal to the values of the triangles it belongs to. However, in case the value of the edge is equal to that of an incident triangle, we make sure we order the edge before the triangle. With this, every prefix of the filter is a complex. Writing K_j for the collection of simplices σ_i with $i \leq j$, we get an increasing sequence of complexes,

$$\emptyset \;=\; K_0 \subseteq K_1 \subseteq \ldots \subseteq K_n \;=\; K. \tag{5.4}$$

We call such an increasing sequence as a *filtration*. The not necessarily contiguous subsequence of alpha complexes is sometimes referred to as the *alpha complex* *filtration* of the set S. This and other filtrations will play an important role in later sections, when we talk about persistent homology.

5.6 Space-Filling Models of Proteins

A major application of alpha complexes are proteins and other molecules modeled as unions of balls [4]. One such model is the *van der Waals diagram* of a protein. It is based on the van der Waals force, which is weakly attractive at short distance between the atoms, and turns into a strongly repulsive force if we push the atoms closer together. The diagram is obtained by taking the union of the balls centered at the atoms in which the radii are chosen so that the atoms are at equilibrium when the balls touch.

Different types of atoms affect neighboring atoms differently, which leads to different radii. For example, hydrogen atoms are the smallest, with carbon, oxygen, and nitrogen atoms represented by somewhat larger balls. This motivates the concept of *weighted alpha complexes*, which are defined analogous to weighted Voronoi diagrams and weighted Delaunay triangulations. To be specific, we have a finite set of sites with real weights, and we recall that the bisector of two sites under the power distance is the set of points x that satisfy $\|x - s\|^2 - w_s = \|x - t\|^2 - w_t$. We have shown that the bisector is a straight line. If we add the same constant, α^2, to the weights, then the bisector stays the same. We therefore define $D_s(\alpha)$ as the disk with center s and radius $\sqrt{w_s + \alpha^2}$. We note that the radius depends on the weight as well as on α. We could therefore drop α and change all weights. We prefer to keep the weights fixed and modify the complex by varying the parameter, thus stressing that we use only one degree of freedom. Observe that the radius in the weighted case agrees with the definition in the unweighted case, when it is α. To construct the weighted α-complex of S, we take the union of the disks $D_s(\alpha)$, for all $s \in S$, we decompose the union into convex regions using the weighted Voronoi diagram, and we take the dual, as before.

If we choose the weights so that the radii agree with the van der Waals forces, we get the weighted 0-complex as the dual of the van der Waals diagram. There are reasons to also consider non-zero values of α. One is the *solvent accessible diagram* in which the van der Waals radii are increased by about 1.4 Angstrom, which is the radius used to approximate a water molecule. The diagram thus represents the interaction between the protein and solvent, which is water. Indeed, if we represent a water molecule by the ball with center x and radius 1.4 Angstrom, then this molecule is not yet repelled by the protein iff x lies outside the solvent accessible model. For other solvents, we would modify the radius by different amounts. A word or caution is in order: increasing all radii by the same amount relates to the Apollonius and not the power diagram. Indeed, only if we increase all squared weights by the same amount—which for reasons of compatibility with the unweighted case is denoted as α^2—we preserve the power diagram.

References

1. Jarvis RA (1973) On the identification of the convex hull of a finite set of points in the plane. Inform Process Lett 2:18–21
2. Jarvis RA (1977) Computing the shape hull of points in the plane. In: Proceedings of IEEE computer society conference pattern recognition and image processing, pp 231–241
3. Edelsbrunner H, Kirkpatrick DG, Seidel R (1983) On the shape of a set of points in the plane. IEEE Trans Inform Theory 29:551–559
4. Lee B, Richards FM (1971) The interpretation of protein structures: estimation of static accessibility. J Mol Biol 55:379–400