

Weighted Alpha Shapes¹

Herbert Edelsbrunner²

Abstract. This paper introduces the concept of an α -shape of a finite set of points with real weights in \mathbb{R}^d . It is a polytope uniquely determined by the points, their weights, and a parameter $\alpha \in \mathbb{R}$ that controls the desired level of detail. The relationship of this concept to regular triangulations and to space filling diagrams is explained. Explicit formulas needed to compute the α -shape from the regular triangulation of the weighted points are given.

Key Words and Phrases. Geometric algorithms, d -dimensional space, point sets, weights, shapes, regular triangulations, power diagrams, space filling diagrams. Combinatorial topology, simplicial complexes, underlying spaces. Computer applications, object representation, surface reconstruction.

¹This work is supported by the Alan T. Waterman award, grant CCR-9118874. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author and do not necessarily reflect the view of the National Science Foundation.

²Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA.

1 Introduction

The α -shape of a finite point set is a polytope that is uniquely determined by the set and a real number α . It expresses the intuitive notion of the shape of the point set, and α is a parameter that controls the level of detail reflected by the polytope. The original paper on α -shapes [9] defines the concept in \mathbb{R}^2 . An extension to \mathbb{R}^3 together with an implementation is reported in [11]. In both papers the relationship between α -shapes and Delaunay triangulations [7] is described in detail and used as the basis of an algorithm for constructing α -shapes.

These algorithms have been implemented by Patrick Moran in \mathbb{R}^2 and by Ernst Mücke in \mathbb{R}^3 , complete with graphic interface. The availability of these implementations, in particular the one in \mathbb{R}^3 , has led to applications in various areas of science and engineering. Some of these applications are briefly described in [11]. A question that was repeatedly asked is whether it is possible to construct a shape that represents different levels of detail in different parts of space. This is not possible with the α -shape as originally defined. In this paper we describe a way to achieve different levels of detail in a single shape by assigning weights to the points. The resulting concept, called the weighted α -shape, is defined and described in real space of arbitrary finite dimension. The unweighted α -shape is the special case obtained by setting all weights equal to zero.

What are the applications where weights can be beneficial?

- (i) In biology and chemistry a common computational task is modeling molecular structures. It is natural to use α -shapes for this purpose as they are precise duals of the popular space filling diagrams obtained by taking the union of balls, one around each atom, see e.g. [15]. Weights are needed to model atoms with different van der Waals radii. This is discussed in section 4.
- (ii) In reconstructing a surface from scattered point data it is rarely the case that the points are uniformly dense everywhere on the (unknown) surface. Indeed, the density often varies with the curvature. If α is chosen so that the α -shape produces a piecewise linear surface in sparse regions, it will be clumsy and hide details in denser regions. Conversely, if α is chosen so that dense regions are nicely modeled then the α -shape will develop holes and break apart in sparse regions. The assignment of large weights in sparse regions and of small weights in dense regions can be used to counteract this undesirable effect.
- (iii) Another goal that can be achieved by assigning weights is to enforce certain edges or faces. These might be given as part of the input, but they cannot be processed directly since α -shapes have been defined only for finite point sets and not for other geometric objects.

Outline. Sections 2 and 3 give two alternative definitions of weighted α -shapes. The approach in section 2 is more direct but it is difficult to see that the definitions are consistent and free of contradictions. Section 3 introduces weighted α -shapes via subcomplexes of so-called regular triangulations, see e.g. [16]. This approach requires some concepts from discrete geometry and algebraic topology. The benefits of this approach are a proof that our definitions are consistent and a method for effectively constructing weighted α -shapes. The duality between α -shapes and so-called space filling diagrams [15] is explained in section 4. The results of section 3 suggest that α -shapes be computed from regular triangulations; efficient algorithms for the latter are known, see e.g. [13]. Section 5 explains how a weighted α -shape can be derived from the regular triangulation of the points. Section 6 concludes this paper with remarks

and open questions. Appendix A contains a tedious proof of an important result stated in section 3, and appendix B provides explicit formulas for the geometric primitives that are needed to construct α -shapes.

2 Weighted Alpha Shapes

Weighted points and orthogonality. As mentioned in the introduction, the weighted α -shape is defined for a finite set of weighted points. Let $S \subseteq \mathbb{R}^d \times \mathbb{R}$ be such a set. A weighted point is denoted as $p = (p', p'')$, with $p' \in \mathbb{R}^d$ its *location* and $p'' \in \mathbb{R}$ its *weight*. For two weighted points, $p = (p', p'')$ and $x = (x', x'')$, define

$$\pi(p, x) = |p'x'|^2 - p'' - x'',$$

where $|p'x'|$ is the Euclidean distance between their locations. Call p and x *orthogonal* if $\pi(p, x) = 0$. The choice of words is motivated by the common interpretation of p as a $(d - 1)$ -sphere with center p' and radius $\sqrt{p''}$: if p'' and x'' are both positive then $\pi(p, x) = 0$ iff the two $(d - 1)$ -spheres intersect at a right angle. If x is unweighted, that is, $x'' = 0$, then $\pi(p, x)$ is the same as the power distance of x' from p , see e.g. [6]. In this case we abuse notation and use $\pi(p, x')$ and $\pi(p, x)$ interchangeably.

For a weighted point p and a real α define $p_{+\alpha} = (p', p'' + \alpha)$. So p and $p_{+\alpha}$ share the same location and their weights differ by α . For example, $p_{-p''} = (p', 0)$ is the unweighted point at the same location as p . The following result is an immediate consequence of the definitions.

2.1 For all $\alpha \in \mathbb{R}$, $\pi(p, x) = \pi(p_{+\alpha}, x_{-\alpha})$.

As a special case of 2.1 we have $\pi(p, x) = \pi(p_{+x''}, x_{-x''})$. This means, for example, that if p and x are orthogonal then $(x_{-x''})'$ lies on the sphere with center p' and radius $\sqrt{p'' + x''}$. We will use the convenient notation $S' = \{p' \mid p \in S\}$ for the projection of S into \mathbb{R}^d .

Definition of weighted alpha shape. For every fixed $\alpha \in \mathbb{R}$, the α -shape of S will be defined as a polytope in \mathbb{R}^d that is uniquely determined by the points in S . A precise definition of what we mean by a polytope can be found in section 3. If all weights are 0 then it will coincide with the (unweighted) α -shape as defined for dimensions 2 and 3 in [9, 11]. In this section, the polytope is specified in terms of its faces. Throughout, we explicitly or implicitly make various general position assumptions. One such assumption is that any $k + 1 \leq d + 1$ points of S' are affinely independent. Another is that for every subset of $d + 1$ weighted points of S there is a unique $x = (x', x'')$ that is orthogonal to all points of this subset, see also section 3. Yet another assumption is that $x'' \neq \alpha$.

Consider a subset $T \subseteq S$, with $|T| = k + 1 \leq d$. It *spans* a k -simplex $\sigma_T = \text{conv}(T')$. We call σ_T *α -exposed* if there exists a weighted point $x = (x', \alpha)$ so that

$$\pi(p, x) \begin{cases} = 0 & \text{for all } p \in T, \text{ and} \\ > 0 & \text{for all } p \in S - T. \end{cases}$$

The *weighted α -shape* of S , $\mathcal{W}_\alpha = \mathcal{W}_\alpha(S)$, is a polytope whose boundary, $\partial\mathcal{W}_\alpha$, is the union of all α -exposed simplices spanned by subsets of S . These simplices are the *faces* of \mathcal{W}_α .

At this moment it is not obvious that the collection of α -exposed faces indeed forms the boundary of a polytope, but this will be discussed in section 3 and formally proved in appendix A. It should be noted

that the polytope is not necessarily connected and that it is not necessarily the same as the closure of its interior.

We next specify which components of $\mathbb{R}^d - \partial\mathcal{W}_\alpha$ are interior to \mathcal{W}_α and which components are exterior. The α -exposed $(d - 1)$ -simplices are the facets of \mathcal{W}_α . Consider such a facet $\sigma = \sigma_T$. There is either one weighted point $x = (x', \alpha)$ that identifies σ as α -exposed, or there are two such points. In the first case, σ bounds the interior of \mathcal{W}_α , namely the side of σ opposite x' belongs to \mathcal{W}_α and the side of x' does not. In the second case, σ does not bound the interior of \mathcal{W}_α . Again it is not obvious that the specification is free of contradictions; that it is will become clear later.

An interesting special case is the weighted α -shape for $\alpha = +\infty$. This is the same as \mathcal{W}_α for sufficiently large α because a finite point set has only finitely many different \mathcal{W}_α . For $T \subseteq S$, σ_T is $(+\infty)$ -exposed iff it is a face of the convex hull of S' . Furthermore, the interior of $\text{conv}(S')$ belongs to $\mathcal{W}_{+\infty}$, so we conclude that $\mathcal{W}_{+\infty} = \text{conv}(S')$.

Varying weights instead of alpha. Because of 2.1 it is also possible to define \mathcal{W}_α in terms of \mathcal{W}_0 for points at the same locations but with different weights. Define $S_{+\alpha} = \{p_{+\alpha} \mid p \in S\}$. For example, $S_{+0} = S$. As a consequence of the definitions and of 2.1 we have the following result.

$$2.2 \text{ For all } \alpha \in \mathbb{R}, \mathcal{W}_\alpha(S) = \mathcal{W}_0(S_{+\alpha}).$$

The collection $\{\mathcal{W}_\alpha(S) \mid \alpha \in \mathbb{R}\}$ will be referred to as the *family* of weighted α -shapes of S . By 2.2 the family is the same as the collection of \mathcal{W}_0 for all $S_{+\alpha}$, $\alpha \in \mathbb{R}$.

3 Weighted Alpha Complexes

The weighted α -shape of S , \mathcal{W}_α , can also be defined as the underlying space of a subcomplex of the regular triangulation of S . This section explains these terms and presents the alternative definition.

Power diagrams and regular triangulations. As usual let $S \subseteq \mathbb{R}^d \times \mathbb{R}$ be a finite set of weighted points. For two points $p, q \in S$ define

$$\chi_{\{p,q\}} = \{x' \in \mathbb{R}^d \mid \pi(p, x') = \pi(q, x')\}$$

and observe that this is a hyperplane orthogonal to the edge connecting p and q . For example, if $p'' = q''$ then $\chi_{\{p,q\}}$ cuts this edge exactly at its midpoint. More generally, for any subset $T \subseteq S$, with $k + 1 = |T| \geq 2$, we define $\chi_T = \bigcap_{p,q \in T} \chi_{\{p,q\}}$, the set of locations $x' \in \mathbb{R}^d$ with equal power distance, $\pi(p, x')$, from all $p \in T$. Because of general position, χ_T is a $(d - k)$ -flat if $1 \leq k \leq d$, and $\chi_T = \emptyset$ if $k \geq d + 1$.

The hyperplane $\chi_{\{p,q\}}$ bounds two closed half-spaces, namely

$$\text{dom}(p, q) = \{x' \in \mathbb{R}^d \mid \pi(p, x') \leq \pi(q, x')\},$$

and $\text{dom}(q, p)$ defined symmetrically. For each $p \in S$ we can thus define its *cell*, $Z_{\{p\}}$, as the intersection of the half-spaces $\text{dom}(p, q)$, for all $q \in S - \{p\}$. This is a convex polyhedron in \mathbb{R}^d and p is called its *generator*. The collection of cells generated by points in S is known as the *power diagram* of S , $\mathcal{P} = \mathcal{P}(S)$, see e.g. [1].

For a subset $T \subseteq S$ we generalize the definition of a cell to $Z_T = \bigcap_{p \in T} Z_{\{p\}}$. For $|T| \geq 2$, Z_T is certainly contained in χ_T , and because the cells $Z_{\{p\}}$ are convex polyhedra, Z_T is also convex. By assumption of general position, Z_T is either empty or a $(d - k)$ -dimensional polyhedron. Consider the special case where T is a subset of $d + 1$ points of S . Then χ_T is a single point which we denote by y'_T . There is a unique weight y''_T so that $y_T = (y'_T, y''_T)$ is orthogonal to all $p \in T$. It should be clear that $y'_T = Z_T$ iff $\pi(q, y_T) > 0$ for all $q \in S - T$. In any case, the points y_T can be used to introduce the important concept of regularity.

A d -simplex $\sigma_T = \text{conv}(T')$ is called *regular* if $\pi(q, y_T) > 0$ for all $q \in S - T$. The collection of regular d -simplices defines the *regular triangulation* of S , denoted $\mathcal{R} = \mathcal{R}(S)$. The notion of a regular triangulation has recently become popular in the area of discrete geometry, see e.g. [16]. The power diagram and the regular triangulation of S are dual to each other in the following sense.

3.1 For every $0 \leq k \leq d$ and every $T \subseteq S$ with $|T| = k + 1$, σ_T is a regular k -simplex iff $Z_T \neq \emptyset$.

For the rest of this paper the regular triangulation is a more important concept than the power diagram. Still, some of the forthcoming arguments may become more intuitive if one visualizes them using the cells of \mathcal{P} and their intersections.

Simplicial complexes. In the language of piecewise-linear topology, see e.g. [14, 18], \mathcal{R} is a simplicial complex with underlying space $\text{conv}(S')$. We begin with some definitions. Call the empty set a (-1) -dimensional simplex and consider it a face of every simplex. A collection \mathcal{K} of simplices in \mathbb{R}^d is a *simplicial complex* provided it satisfies the following two conditions: (i) every face of a simplex $\sigma \in \mathcal{K}$ is also in \mathcal{K} , and (ii) the intersection of any two simplices in \mathcal{K} is a face of each. So in order to make \mathcal{R} a genuine complex it ought to be defined as the collection of all regular d -simplices *and* their faces. Hence (i) is satisfied.

Why does (ii) hold for \mathcal{R} ? Map each point $p \in S$ to the point $p^+ \in \mathbb{R}^{d+1}$ as follows. Let $p = (p', p'')$ with $p' = (\phi_1, \phi_2, \dots, \phi_d)$, and define $p^+ = (\phi_1, \phi_2, \dots, \phi_d, \phi_{d+1})$, where $\phi_{d+1} = \sum_{i=1}^d \phi_i^2 - p''$. The convex hull of $S^+ = \{p^+ \mid p \in S\}$ is a convex polytope \mathcal{Q} in \mathbb{R}^{d+1} . It is simplicial because we assume general position of the weighted points. Refer to the direction of the $(d + 1)$ st coordinate-axis as *vertical*. Let $T = \{p_0, p_1, \dots, p_d\}$ be a subset of $d + 1$ points of S , and consider the d -simplex $\sigma_T = \text{conv}(T')$. It is fairly easy to see that σ_T is a regular d -simplex of S iff $\text{conv}(T^+)$ is a facet of \mathcal{Q} so that \mathcal{Q} lies vertically above the hyperplane $\text{aff}(T^+)$. This hyperplane is given by the equation

$$\det \begin{pmatrix} 1 & \phi_{0,1} & \phi_{0,2} & \dots & \phi_{0,d+1} \\ 1 & \phi_{1,1} & \phi_{1,2} & \dots & \phi_{1,d+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \phi_{d,1} & \phi_{d,2} & \dots & \phi_{d,d+1} \\ 1 & x_1 & x_2 & \dots & x_{d+1} \end{pmatrix} = 0,$$

see also appendix B. Call such a facet a *lower facet* of \mathcal{Q} . So \mathcal{R} can be obtained by vertically projecting all lower facets of \mathcal{Q} into \mathbb{R}^d , which is spanned by the first d coordinate-axes of \mathbb{R}^{d+1} . Condition (ii) follows for \mathcal{R} because the projection is one-to-one and because (ii) holds for the facets of \mathcal{Q} and their lower-dimensional faces.

Subcomplexes. It follows fairly straightforwardly from the definitions that if σ_T is an α -exposed simplex of S then $\sigma_T \in \mathcal{R}$. In fact, for every k -simplex $\sigma \in \mathcal{R}$, $0 \leq k \leq d - 1$, there is an $\alpha \in \mathbb{R}$ so that σ is

α -exposed. We shed more light on the relationship between the regular triangulation and the α -shapes of S . We do this by considering subcomplexes of the regular triangulation. A subset $\mathcal{L} \subseteq \mathcal{K}$ is a *subcomplex* of \mathcal{K} if it is a complex itself, that is, if (i) also holds for \mathcal{L} .

For every simplex σ_T let $y_T = (y'_T, y''_T)$ be the point with minimum weight that is orthogonal to all $p \in T$. As mentioned before, if $|T| = d + 1$ there is only one weighted point orthogonal to all $p \in T$. If $|T| = k + 1 \leq d$ there are uncountably many such weighted points, but only one minimizes the weight. In fact, the set of locations of these weighted points is the $(d - k)$ -flat χ_T and the one with minimum weight is located at the intersection of χ_T with the orthogonal k -flat $\text{aff}(T')$. Call $\varrho_T = y''_T$ the *size* of σ_T . For example, the size of a point p is $\varrho_{\{p\}} = -p''$. Include $\varrho_T \neq \alpha$ and $\pi(q, y_T) \neq 0$ for all $q \in S - T$ in the collection of general position assumptions. The size of simplices obeys the following monotonicity property.

3.2 If σ_T is a proper face of σ_U then $\varrho_T < \varrho_U$.

A point $q \in S - T$ is a *conflict* for y_T if $\pi(q, y_T) < 0$ and y_T is *conflict-free* if it has no conflict. Incidentally, y_T is conflict-free iff $Z_T \cap \text{aff}(T') \neq \emptyset$. For example, y_T is always conflict-free if σ_T is a d -simplex in \mathcal{R} , but it can have conflicts if it is a simplex in \mathcal{R} whose dimension is less than d . The *weighted α -complex* of S , $\mathcal{K}_\alpha = \mathcal{K}_\alpha(S)$, is the subcomplex of \mathcal{R} that contains $\sigma_T \in \mathcal{R}$ if

(C1) $\varrho_T < \alpha$ and y_T is conflict-free, or

(C2) σ_T is a face of another simplex in \mathcal{K}_α .

A simplex in \mathcal{K}_α is a *principal simplex* if it is not a proper face of any other simplex in \mathcal{K}_α . Hence, all principal simplices satisfy (C1), but it is possible that a simplex that satisfies (C1) is not principal.

Underlying spaces. The *underlying space* of a complex \mathcal{K} is $|\mathcal{K}| = \bigcup_{\sigma \in \mathcal{K}} \sigma$. For example, $|\mathcal{R}| = \text{conv}(S')$. In this paper, and more generally in the algebraic topology literature, a *polytope* is defined as the underlying space of a finite simplicial complex. The most important result on weighted α -complexes is that they provide an alternative way to construct the polytopes called weighted α -shapes in section 2. Specifically, we have the following result.

3.3 For all $\alpha \in \mathbb{R}$, $\mathcal{W}_\alpha = |\mathcal{K}_\alpha|$.

The proof of 3.3 is fairly tedious and can be found in appendix A of this paper.

If σ_T is a d -simplex of \mathcal{R} then y_T is conflict-free by definition of \mathcal{R} . Hence, $\sigma_T \in \mathcal{K}_\alpha$ iff $\varrho_T < \alpha$. Let now σ_T be a k -simplex of \mathcal{R} , with $k < d$. If y_T is conflict-free then again we have $\varrho_T < \alpha$ as the only condition for $\sigma_T \in \mathcal{K}_\alpha$. Otherwise, σ_T is a face of another simplex in \mathcal{K}_α . This implies the following monotonicity property for weighted α -shapes and α -complexes.

3.4 For all $\alpha_1 \leq \alpha_2$, \mathcal{K}_{α_1} is a subcomplex of \mathcal{K}_{α_2} and therefore $\mathcal{W}_{\alpha_1} \subseteq \mathcal{W}_{\alpha_2}$.

Section 5 will pick up on the above idea of determining thresholds for each simplex that help decide for which values of α it belongs to \mathcal{K}_α . Based on these thresholds and the corresponding formulas in appendix B we will be able to quickly identify \mathcal{K}_α . Weighted α -shapes can thus be constructed by first computing \mathcal{R} and then selecting the desired simplices. The construction of \mathcal{R} is the algorithmically more demanding step and we refer to [13] for details.

4 Space Filling Diagrams

A popular concept in computational biology and chemistry is the so-called space filling diagram, see [15, 19, 5]. It models a molecule by representing its atoms as overlapping balls in \mathbb{R}^3 . The radius of a ball is chosen so that its bounding sphere approximates the van der Waals surface of the atom. Different ball sizes are necessary for different types of atoms. This is where weights are needed.

Definition of space filling diagram. For a weighted point $p = (p', p'') \in \mathbb{R}^d \times \mathbb{R}$ define the closed d -ball $B_p = \{x' \in \mathbb{R}^d \mid |x'p'|^2 \leq p''\}$. If $p'' < 0$ then $B_p = \emptyset$. The *space filling diagram* of $S \subseteq \mathbb{R}^d \times \mathbb{R}$ is

$$\mathcal{F}_0 = \mathcal{F}_0(S) = \bigcup_{p \in S} B_p.$$

A point $p \in S$ is *redundant* if $\mathcal{F}_0(S) = \mathcal{F}_0(S - \{p\})$. For example, all p with negative weight are redundant. For general $\alpha \in \mathbb{R}$, we define $\mathcal{F}_\alpha = \mathcal{F}_\alpha(S) = \mathcal{F}_0(S_{+\alpha})$. Intuitively, \mathcal{F}_α can be derived from \mathcal{F}_0 by blowing up or shrinking the balls. Note that the change affects the balls in a non-linear manner. More precisely, the squares of the radii all grow or shrink by the same amount, namely α .

It follows from 2.1 that \mathcal{F}_α consists of all locations x' of weighted points $x = (x', \alpha)$ so that $\pi(p, x) \leq 0$ for at least one $p \in S$. Furthermore, $x' \in \partial\mathcal{F}_\alpha$ iff $\pi(p, x) \geq 0$ for all $p \in S$ and $\pi(p, x) = 0$ for at least one $p \in S$. For any arbitrary point $x = (x', \alpha)$ let T_x be the subset of points $p \in S$ for which $\pi(p, x) = 0$. The following property follows from the definitions.

4.1 $x' \in \partial\mathcal{F}_\alpha$ iff σ_{T_x} is α -exposed.

Details of the relationship between \mathcal{F}_α and \mathcal{W}_α indicated by 4.1 will be explained shortly.

Why the square of the radius? We explain why we increase the *weight* of the points by a uniform amount in the above definition of the family \mathcal{F}_α . Natural alternatives would be to increase the radii of the B_p by α or to multiply the radii by α . Consider a single d -ball B_p under the proposed model. As α increases continuously from $-\infty$ to $+\infty$, the contribution of ∂B_p to $\partial\mathcal{F}_\alpha$ is a changing portion of a growing $(d-1)$ -sphere. It is fairly easy to see that this portion sweeps out a convex polyhedron, namely the cell $Z_{\{p\}}$ in the power diagram $\mathcal{P} = \mathcal{P}(S)$, see section 3. Recall that \mathcal{P} is the dual of the regular triangulation, \mathcal{R} . Now repeat the same experiment but increase the *radii* of the d -balls by α . The portion of ∂B_p on $\partial\mathcal{F}_\alpha$ sweeps out a region in \mathbb{R}^d that is connected but not necessarily convex. This region is the cell of p in the additively weighted Voronoi diagram of S , see e.g. [17]. If the radii of the balls are multiplied by α then the region swept out by the portion of ∂B_p on $\partial\mathcal{F}_\alpha$ is no longer necessarily connected. In this model it is the (possibly disconnected) cell of p in the multiplicatively weighted Voronoi diagram of S , see e.g. [3]. Of the three diagrams, the power diagram is computationally most tractable and it is the only one with a natural dual, namely the regular triangulation.

Duality between \mathcal{F}_α and \mathcal{W}_α . The boundary of \mathcal{F}_α consists of pieces of spheres of various dimensions. More specifically, let $T \subseteq S_{+\alpha}$, $|T| = k+1 \leq d$, so that $\sigma_T \in \mathcal{R}$ and $\bigcap_{p \in T} B_p \neq \emptyset$. Because of the general position assumption, the intersection of the corresponding $(d-1)$ -spheres,

$$K_T = \bigcap_{p \in T} \partial B_p,$$

is an ℓ -sphere, with $\ell = d - k - 1$. For example, if $|T| = d$ then K_T is a pair of points, and if $|T| = d - 1$ then K_T is a circle. An ℓ -face of \mathcal{F}_α is a connected component of the intersection of $\partial\mathcal{F}_\alpha$ with an ℓ -sphere

K_T . If T is a subset of $S_{+\alpha}$ so that the ℓ -sphere K_T , $\ell = d - |T|$, contains at least one ℓ -face of \mathcal{F}_α then there is a weighted point $x = (x', \alpha)$ with $T_x = T$. By 4.1, σ_T is therefore α -exposed and thus a face of \mathcal{W}_α . The argument can also be made in the other direction. We therefore have the following result.

4.2 Let $T \subseteq S_{+\alpha}$ with $0 \leq |T| = k + 1 \leq d$ and set $\ell = d - k - 1$. Then σ_T is a k -face of \mathcal{W}_α iff K_T contains at least one ℓ -face of \mathcal{F}_α .

Given the faces of \mathcal{W}_α it is thus straightforward to identify the spheres that contribute faces to the boundary of \mathcal{F}_α . By exploiting the incidences between the faces of \mathcal{W}_α one can also recover the faces of \mathcal{F}_α and their incidence structure. This will be described elsewhere.

5 Face Classification

There are simple necessary and sufficient conditions when a k -simplex $\sigma_T \in \mathcal{R}$ belongs to \mathcal{K}_α that can be derived from the definition of \mathcal{K}_α in section 3. We refine these conditions to distinguish between three types of simplices in \mathcal{K}_α depending on their relationship to \mathcal{W}_α . Recall that a principal simplex of \mathcal{K}_α is not a proper face of any other simplex in \mathcal{K}_α . Call a k -simplex $\sigma_T \in \mathcal{K}_\alpha$

singular if σ_T is a face of \mathcal{W}_α and also a principal simplex of \mathcal{K}_α ,
regular if σ_T is a face of \mathcal{W}_α and not a principal simplex of \mathcal{K}_α , and
interior if it is not a face of \mathcal{W}_α .

Note that every d -simplex in \mathcal{K}_α is interior by definition. This classification is somewhat arbitrary and a finer differentiation, in particular of the set of regular simplices, can be established if need be.

The classification of σ_T as singular, regular, or interior is related to a certain subcomplex of \mathcal{K}_α defined for σ_T . The *link* of $\sigma_T \in \mathcal{K}_\alpha$ is

$$\text{Lk}_\alpha(\sigma_T) = \{\sigma_V \in \mathcal{K}_\alpha \mid T \cap V = \emptyset, \sigma_{T \cup V} \in \mathcal{K}_\alpha\}.$$

As an example consider the link of σ_T in $\mathcal{R} = \mathcal{K}_\infty$. We call a complex \mathcal{L} a *topological i -sphere* if its underlying space is homeomorphic to a geometric i -sphere, and we call it a *topological i -ball* if it is homeomorphic to a geometric i -ball. If a k -simplex σ_T is a face of $\text{conv}(S')$ then $\text{Lk}_\infty(\sigma_T)$ is a topological ℓ -ball, where $\ell = d - k - 1$. On the other hand, if σ_T does not lie on the boundary of $\text{conv}(S')$ then $\text{Lk}_\infty(\sigma_T)$ is a topological ℓ -sphere. The following relationship between the type and the link of σ_T is fairly straightforward. Call the empty set a topological (-1) -sphere.

$$5.1 \quad \sigma_T \in \mathcal{K}_\alpha \text{ is } \begin{cases} \text{singular} & \text{if } \text{Lk}_\alpha(\sigma_T) = \emptyset \text{ and } \ell \geq 0, \\ \text{interior} & \text{if } \text{Lk}_\alpha(\sigma_T) \text{ is a topological } \ell\text{-sphere, and} \\ \text{regular} & \text{otherwise.} \end{cases}$$

We need some more definitions to express the classification in terms of intervals for α . A k -simplex $\sigma_T \in \mathcal{R}$ is *attached* if there is no $\alpha \in \mathbb{R}$ so that σ_T is a principal simplex of \mathcal{K}_α . Otherwise, σ_T is *unattached*. Notice that all d -simplices are unattached. This notion is related to the point y_T defined in section 3. Indeed, the following result can easily be proved using 3.2.

5.2 σ_T is unattached iff y_T is conflict-free.

For a k -simplex $\sigma_T \in \mathcal{R}$, with $k \leq d - 1$, let $\text{up}(\sigma_T)$ be the set of all simplices that contain σ_T as a proper face, that is, $\text{up}(\sigma_T) = \{\sigma_U \in \mathcal{R} \mid T \subset U\}$. If σ_T is unattached then it belongs to \mathcal{K}_α iff $\varrho_T < \alpha$. Otherwise, $\sigma_T \in \mathcal{K}_\alpha$ iff at least one $\sigma_U \in \text{up}(\sigma_T)$ belongs to \mathcal{K}_α . We thus define

$$\underline{\mu}_T = \min\{\varrho_U \mid \sigma_U \in \text{up}(\sigma_T), \sigma_U \text{ is unattached}\}.$$

If σ_T is attached then $\sigma_T \in \mathcal{K}_\alpha$ iff $\underline{\mu}_T < \alpha$. If σ_T is unattached then it is a non-singular simplex of \mathcal{K}_α iff $\underline{\mu}_T < \alpha$. To distinguish between regular and interior simplices we define

$$\overline{\mu}_T = \max\{\varrho_U \mid \sigma_U \in \text{up}(\sigma_T)\}.$$

By 3.2, $\overline{\mu}_T$ is always the size of a d -simplex. Unless σ_T is a face of $\text{conv}(S')$ it is an interior simplex of \mathcal{K}_α iff $\overline{\mu}_T < \alpha$. Faces of $\text{conv}(S')$ cannot ever be interior simplices. Finally, we set $\underline{\mu}_T = \overline{\mu}_T = \varrho_T$ for all d -simplices σ_T of \mathcal{R} .

Depending on whether or not σ_T is attached and whether or not it is a face of $\text{conv}(S')$ we can thus specify the intervals for α in which σ_T is singular, regular, and interior, see table 5.1. If σ_T is a d -simplex

σ_T	singular	regular	interior
unattached, not on $\partial\text{conv}(S')$	$(\varrho_T, \underline{\mu}_T)$	$(\underline{\mu}_T, \overline{\mu}_T)$	$(\overline{\mu}_T, \infty]$
attached, not on $\partial\text{conv}(S')$	\emptyset	$(\underline{\mu}_T, \overline{\mu}_T)$	$(\overline{\mu}_T, \infty]$
unattached, on $\partial\text{conv}(S')$	$(\varrho_T, \underline{\mu}_T)$	$(\underline{\mu}_T, \infty]$	\emptyset
attached, on $\partial\text{conv}(S')$	\emptyset	$(\underline{\mu}_T, \infty]$	\emptyset

Table 5.1: Intervals of α for which $\sigma_T \in \mathcal{R}$ belongs to \mathcal{K}_α .

of \mathcal{R} then the first row of the table applies. In this case the intervals for the singular and regular columns are empty so that the only interval is $(\overline{\mu}_T, \infty] = (\varrho_T, \infty]$ in the interior column.

Computing intervals. In order to classify a simplex $\sigma_T \in \mathcal{R}$ we need a test that decides whether σ_T is attached or unattached. By the regularity of \mathcal{R} it suffices to check points $q \in S - T$ for which $\sigma_{T \cup \{q\}} \in \mathcal{R}$. We have y_T conflict-free iff no such point q is a conflict of y_T . By 5.2, σ_T is unattached iff $\pi(q, y_T) > 0$ for all such points q .

The intervals can conveniently be computed in the order of decreasing dimension. At the time a k -simplex σ_T is processed the values $\varrho_U, \underline{\mu}_U, \overline{\mu}_U$ for all $(k + 1)$ -simplices $\sigma_U \in \text{up}(\sigma_T)$ are available. Denote this set of $(k + 1)$ -simplices by $\text{up}_{k+1}(\sigma_T)$. The following simple rules are easily verified.

5.3 Let σ_T be a k -simplex of \mathcal{R} , with $0 \leq k \leq d - 1$.

- (i) $\underline{\mu}_T$ is the minimum of the ϱ_U , over all unattached $\sigma_U \in \text{up}_{k+1}(\sigma_T)$, and the $\underline{\mu}_U$, over all attached $\sigma_U \in \text{up}_{k+1}(\sigma_T)$.
- (ii) $\overline{\mu}_T$ is the maximum of the $\overline{\mu}_U$, over all $\sigma_U \in \text{up}_{k+1}(\sigma_T)$.

We assess the time needed by an algorithm that computes the values $\varrho, \underline{\mu}$, and $\overline{\mu}$ using 5.3. The assumption is that d is a constant and that for a given σ_T the $(k + 1)$ -simplices in $\text{up}_{k+1}(\sigma_T)$ can be accessed in constant time per simplex. The time required for σ_T is therefore proportional to the cardinality of

$\text{up}_{k+1}(\sigma_T)$. The sum of cardinalities, over all simplices in \mathcal{R} , is the same as the sum of $|T| + 1$, over all simplices $\sigma_T \in \mathcal{R}$, which is less than $d + 1$ times the total number of simplices in \mathcal{R} . We conclude that the running time is proportional to the size of the set \mathcal{R} .

6 Remarks

The main contribution of this paper is the generalization of α -shapes [9, 11] to a weighted environment. The new concept, the weighted α -shape, is closely related to regular triangulations and to space filling diagrams. All definitions and results are given for real space of arbitrary finite dimension. Weighted α -shapes will be implemented and applications including the ones mentioned in the introduction will be explored. It is planned to report on these experiments later and elsewhere.

A fascinating aspect of weighted α -shapes is the possibility to influence the generation of shapes through different weight assignments. On the one hand, this makes a very flexible tool with a great facility for adaptation. On the other hand, the amount of flexibility offered by free weight assignment is currently not well understood. Here are two sample problems worth studying.

- (i) Given a finite point set in \mathbb{R}^3 , decide whether there is an assignment of real weights so that \mathcal{W}_0 is a closed surface or a collection of disjoint closed surfaces.

A question related to the two-dimensional version of problem (i) has been studied in [12]. The goal is to find weights so that \mathcal{W}_0 is a closed cycle or a collection of disjoint closed cycles. The result is that for n points in \mathbb{R}^2 it is possible to find such weights in time $O(n^2 \log n)$, if they exist.

- (ii) Let S be a set of n points in \mathbb{R}^3 and let E be a collection of non-crossing edges and triangles defined by S . How fast can one decide whether there are weights for the points so that all edges and triangles in E belong to the thus defined regular triangulation of S ?

The special case of deciding whether or not a given triangulation is regular has been studied by several authors, see e.g. [2, 21]. A related problem is to decide whether there is *any* triangulation of S that contains E . This is known to be NP-complete [20]. Maybe the study of secondary polytopes, see e.g. [4], can shed some light on these and related problems.

Acknowledgements

I thank Tamal Dey, Edgar Ramos, and Nimish Shah for helpful comments on an earlier version of this paper.

References

- [1] F. Aurenhammer. Power diagrams: properties, algorithms, and applications. *SIAM J. Comput.* **16** (1987), 78–96.

- [2] F. Aurenhammer. A criterion for the affine equivalence of cell complexes in \mathbb{R}^d and convex polyhedra in \mathbb{R}^{d+1} . *Discrete Comput. Geom.* **2** (1987), 49–64.
- [3] F. Aurenhammer and H. Edelsbrunner. An optimal algorithm for constructing the weighted Voronoi diagram in the plane. *Pattern Recognition* **17** (1984), 251–257.
- [4] L. J. Billera, P. Filliman and B. Sturmfels. Constructions and complexity of secondary polytopes. *Adv. Math.* **83** (1990), 155–179.
- [5] M. L. Connolly. Analytical molecular surface calculation. *J. Appl. Cryst.* **16** (1983), 548–558.
- [6] H. S. M. Coxeter and S. L. Greitzer. *Geometry Revisited*. Math. Assoc. Amer., Washington, DC, 1967.
- [7] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvenyka Nauk* **7** (1934), 793–800.
- [8] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Heidelberg, Germany, 1987.
- [9] H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551–559.
- [10] H. Edelsbrunner and E. P. Mücke. Simulation of Simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Comput. Graphics* **9** (1990), 66–104.
- [11] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. In “Proc. Workshop on Volume Visualization, 1992”, 75–82.
- [12] H. Edelsbrunner, G. Rote and E. Welzl. Testing the necklace condition for shortest tours and optimal factors in the plane. *Theoret. Comput. Sci.* **66** (1989), 157–180.
- [13] H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. In “Proc. 8th Ann. Sympos. Comput. Geom., 1992”, 43–52.
- [14] P. J. Giblin. *Graphs, Surfaces, and Homology*. Second edition, Chapman and Hall, London, 1981.
- [15] B. Lee and F. M. Richards. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.
- [16] C. Lee. Regular triangulations of convex polytopes. In *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, 443–456, P. Gritzmann and B. Sturmfels, eds., American Mathematical Society, Providence, RI, 1991.
- [17] D. T. Lee and R. L. Drysdale III. Generalization of Voronoi diagrams in the plane. *SIAM J. Comput.* **10** (1981), 73–87.
- [18] J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, California, 1984.
- [19] F. M. Richards. Areas, volumes, packing, and protein structure. *Ann. Rev. Biophys. Bioeng.* **6** (1977), 151–176.
- [20] J. Ruppert and R. Seidel. On the difficulty of triangulating three-dimensional nonconvex polyhedra. *Discrete Comput. Geom.* **7** (1992), 227–253.
- [21] F. Supnick. On the perspective deformation of polyhedra. *Ann. Math.* **49** (1948), 714–730, and **53** (1953), 551–555.

Appendix A Proof of 3.3

This appendix presents a proof of the important result 3.3 stated in section 3. At several places in the proof we use a point $v = (v', v'')$ that continuously moves while remaining orthogonal to all points of some set $T \subseteq S$. We first discuss the constraints on this continuous motion imposed by T .

Let $T = \{p_0, p_1, \dots, p_k\}$ be a set of $k + 1$ weighted points in $\mathbb{R}^d \times \mathbb{R}$, for some $0 \leq k \leq d$. Recall that v is orthogonal to p_i , $0 \leq i \leq k$, iff

$$\pi(p_i, v) = |p'_i v'|^2 - p''_i - v'' = 0.$$

This relation by itself does not constrain the location of v , but for a given v' there is only one choice for v'' . For two indices, $0 \leq i < j \leq k$, the relations $\pi(p_i, v) = \pi(p_j, v) = 0$ constrain v' to a hyperplane normal to the edge connecting p_i and p_j . This is the hyperplane $\chi_{\{p_i, p_j\}}$ of section 3. The $k + 1$ points of T give rise to a total of $k + 1$ relations constraining v' to the $(d - k)$ -flat χ_T defined in section 3.

The location v' of v can move freely in χ_T , and v'' can always be adjusted so that $\pi(p_i, v) = 0$ for all $0 \leq i \leq k$. Conversely, if $d - k \geq 1$ it is possible to freely increase v'' and to adjust v' in χ_T so that v remains orthogonal to all $p_i \in T$. Recall that $y_T = (y'_T, y''_T)$ is the point with minimum weight orthogonal to all points in T . So $y'_T \in \chi_T$, and indeed $y'_T = \chi_T \cap \text{aff}(T')$, as mentioned in section 3. At a certain step in the proof of 3.3, v' will move continuously while v'' remains fixed and v remains orthogonal to all $p_i \in T$. The set of possible locations is a $(d - k - 1)$ -sphere with center y'_T contained in χ_T . Provided $d - k - 1 \geq 1$ this set is connected and every point can be reached by a continuous motion within the set.

The remainder of this appendix restates 3.3 and presents a proof of this result. It is useful to recall the definitions of \mathcal{W}_α and \mathcal{K}_α before reading the proof.

3.3 For all $\alpha \in \mathbb{R}$, $\mathcal{W}_\alpha = \|\mathcal{K}_\alpha\|$.

Proof. Fix $\alpha \in \mathbb{R}$ and consider values of k between 0 and $d - 1$ inclusive. First we prove that all k -faces of \mathcal{W}_α are contained in the boundary of $\|\mathcal{K}_\alpha\|$, second we show that all k -simplices of \mathcal{K}_α contained in the boundary of $\|\mathcal{K}_\alpha\|$ are faces of \mathcal{W}_α , and third we consider the interiors of \mathcal{W}_α and $\|\mathcal{K}_\alpha\|$.

For the first step let σ_T be a k -face of \mathcal{W}_α . So σ_T is α -exposed which means there exists a weighted point $x = (x', \alpha)$ with $\pi(p, x) = 0$, for all $p \in T$, and $\pi(q, x) > 0$, for all $q \in S - T$. Set $v = x$, increase the weight of v continuously, and adjust its location so that π vanishes for all $p \in T$ and remains non-negative for all $q \in S - T$. Eventually, the weight of v exceeds every finite bound, in which case σ_T is a face of $\text{conv}(S')$, or v reaches a point y_U , $T \subseteq U$, of some d -simplex σ_U . In the latter case we have $\varrho_U > \alpha$ so $\sigma_U \notin \mathcal{K}_\alpha$. In either case it is true that if $\sigma_T \in \mathcal{K}_\alpha$ then it is contained in the boundary of $\|\mathcal{K}_\alpha\|$. Because of general position the existence of x implies that y_T has weight smaller than α , so $\varrho_T = y''_T < \alpha$. If y_T is conflict-free then $\sigma_T \in \mathcal{K}_\alpha$ by condition (C1) of section 3. Otherwise, there are points $q \in S - T$ with $\pi(q, y_T) < 0$. If σ_T is a $(d - 1)$ -face then there is such a point q so that σ_V , with $V = T \cup \{q\}$, is a d -simplex of \mathcal{R} . Move v' continuously on the edge from x' to y'_T and decrease v'' accordingly. At some moment we have $\pi(q, v) = 0$ which implies that $\varrho_V < \alpha$. It follows that $\sigma_V \in \mathcal{K}_\alpha$ and hence $\sigma_T \in \mathcal{K}_\alpha$ by condition (C2). For $k < d - 1$ we assume inductively that all $(k + 1)$ -faces of \mathcal{W}_α are $(k + 1)$ -simplices in \mathcal{K}_α . Again set $v = x$. Now move v continuously so that its weight remains unchanged and $\pi(p, v) = 0$ for all $p \in T$. For each $q \in S - T$ there is such a v so that $\pi(q, v) < \pi(q, y_T)$. This is because $\chi_{\{v, y_T\}}$ contains all points of T and is therefore a hyperplane that passes through y_T . As v moves this hyperplane pivots

about y_T . In particular, there is such a v for each conflict q of y_T . This implies that there is a point v with $\pi(q, v)$ equal to zero for one point $q \in S - T$ and positive for all others. Stop the motion at such a point v . Set $W = T \cup \{q\}$ and notice that σ_W is an α -exposed $(k + 1)$ -simplex of \mathcal{R} . By induction hypothesis $\sigma_W \in \mathcal{K}_\alpha$ and by condition (C2) we have $\sigma_T \in \mathcal{K}_\alpha$. As mentioned earlier, σ_T is therefore contained in the boundary of $|\mathcal{K}_\alpha|$.

We begin the second step by considering the principal simplices of \mathcal{K}_α . Let σ_T be such a principal k -simplex. By definition of \mathcal{K}_α , y_T is conflict-free and $\varrho_T < \alpha$. If $k < d$ then σ_T is contained in the boundary of $|\mathcal{K}_\alpha|$. Let σ_U be any incident d -simplex in \mathcal{R} . Since σ_U is not in \mathcal{K}_α we have $\varrho_U > \alpha$. Now move a point v' continuously from y'_T to y'_U and define v'' so that $\pi(p, v) = 0$ for all $p \in T$. Stop when $v'' = \alpha$ which must happen between y'_T and y'_U because $y''_T < \alpha < y''_U$. We have $\pi(p, v) > 0$ for all $p \in S - T$ so σ_T is a face of \mathcal{W}_α . Next, consider a non-principal k -face $\sigma_T \in \mathcal{K}_\alpha$ contained in the boundary of $|\mathcal{K}_\alpha|$. If $k = d - 1$ then one incident d -simplex σ_V has $\varrho_V < \alpha$. Assume first that σ_T is not a face of $\text{conv}(S')$. Therefore the other incident d -simplex $\sigma_U \in \mathcal{R}$ has size $\varrho_U > \alpha$. Move v' continuously from y'_V to y'_U and adjust the weight v'' so that $\pi(p, v) = 0$ for all $p \in T = U \cap V$. Again we stop the motion when $v'' = \alpha$ which must occur because $y''_V < \alpha < y''_U$. As before, σ_T is therefore a face of \mathcal{W}_α . If σ_T is a face of $\text{conv}(S')$ and σ_U therefore does not exist then the argument is similar. The only change is that now v' moves continuously from y'_V to infinity along a path so that $\pi(q, v) > 0$ for all $q \in S - T$. Finally, let $k < d - 1$ and assume inductively that all $(k + 1)$ -simplices of \mathcal{K}_α contained in the boundary of $|\mathcal{K}_\alpha|$ are faces of \mathcal{W}_α . Since σ_T is not principal it is a face of some such $(k + 1)$ -simplex σ_U . So σ_U is a $(k + 1)$ -face of \mathcal{W}_α which implies that σ_T is a k -face of \mathcal{W}_α .

The third step shows that the interior of \mathcal{W}_α is the same as the interior of $|\mathcal{K}_\alpha|$. We already proved that $\partial\mathcal{W}_\alpha = \partial|\mathcal{K}_\alpha|$. This boundary decomposes \mathbb{R}^d into a number of connected components. Since these components are the same for \mathcal{W}_α and $|\mathcal{K}_\alpha|$, we just need to show that a component belongs to \mathcal{W}_α iff it contains at least one d -simplex of \mathcal{K}_α . The unbounded component is, of course, exterior to both. Each bounded component, C , is the connected interior of the union of some d -simplices in \mathcal{R} . Let σ_T be a facet of C and let σ_U be the d -simplex with facet σ_T whose interior lies in C . Let $\{r\} = U - T$. Now, $\sigma_U \in \mathcal{K}_\alpha$ iff $\varrho_U < \alpha$ because y_U is conflict-free by definition of \mathcal{R} . Since σ_T is a $(d - 1)$ -simplex there are two points with weight α that are orthogonal to all $p \in T$. Let $x = (x', \alpha)$ be the one so that x' lies on the same side of $\text{aff}(T')$ as σ_U . We have $\pi(q, x) > 0$ for all $q \in S - T$ iff $\pi(r, x) > 0$. So $\sigma_U \notin \mathcal{K}_\alpha$ iff x is a witness of σ_T being α -exposed. The latter is exactly the condition that C is exterior to \mathcal{W}_α . \square

Appendix B Geometric Primitives

The construction of the regular triangulation of a set S of n weighted points in \mathbb{R}^d requires only two types of geometric tests. We call the first a hyperplane test; it decides on which side of a hyperplane spanned by d given points a $(d + 1)$ st given point lies. The second is the orthogonality test; it decides the sign of π measured between a $(d + 2)$ nd point and the unique weighted point orthogonal to $d + 1$ given points. Once \mathcal{R} is constructed, we need to compute the intervals as described in section 5. For this, the size of each simplex $\sigma_T \in \mathcal{R}$ must be computed. Furthermore, it is necessary to test whether σ_T is attached or unattached. This appendix derives an explicit formula for each test. Implementation decisions, such as how to evaluate a determinant, whether or not to use integer arithmetic, and how to resolve ambiguities arising from degenerate data, are left to the programmer. We will, however, assume there is no degeneracy in our data, which can be simulated by programming techniques such as SoS [10].

Let $S = \{p_1, p_2, \dots, p_n\}$, where $p_i = (p'_i, p''_i)$ and $p'_i = (\phi_{i,1}, \phi_{i,2}, \dots, \phi_{i,d})$. It will be convenient to set $\phi_{i,0} = 1$ and $\phi_{i,d+1} = \sum_{j=1}^d \phi_{i,j}^2 - p''_i$, for all $1 \leq i \leq n$. We introduce the following auxiliary points in \mathbb{R}^{d+1} . For $-(d+1) \leq -i \leq -1$ denote by p_{-i} the point at infinity with homogeneous coordinates $(\phi_{-i,0}; \phi_{-i,1}, \phi_{-i,2}, \dots, \phi_{-i,d+1})$, where $\phi_{-i,i} = 1$ and $\phi_{-i,j} = 0$ for $j \neq i$. The special index x is used for the generic point $p_x = (x_0 = 1; x_1, x_2, \dots, x_{d+1})$. The geometric primitives will be expressed using minors with the notation

$$\mathcal{M}_{j_1, j_2, \dots, j_k}^{i_1, i_2, \dots, i_k} = \det \begin{pmatrix} \phi_{i_1, j_1} & \phi_{i_1, j_2} & \cdots & \phi_{i_1, j_k} \\ \phi_{i_2, j_1} & \phi_{i_2, j_2} & \cdots & \phi_{i_2, j_k} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{i_k, j_1} & \phi_{i_k, j_2} & \cdots & \phi_{i_k, j_k} \end{pmatrix}.$$

Hyperplane test. The input to this test consists of $d+1$ points given by their indices i_0, i_1, \dots, i_d between 1 and n inclusive. The first d points span a hyperplane, h , and their sequence imposes an orientation on h . We call one open half-space defined by h its *positive* and the other its *negative side*. The two sides are distinguished by evaluating the sign of the following minor.

B.1 Point p_{i_d} lies on the positive side of h iff $\mathcal{M}_{0,1,\dots,d}^{i_0,i_1,\dots,i_d} > 0$.

Of course, p_{i_d} lies on the negative side of h iff the minor is negative. The case where the minor vanishes corresponds to the degenerate case where p_{i_d} lies on h .

Orthogonality test. The input are $d+2$ points with indices i_0, i_1, \dots, i_{d+1} between 1 and n inclusive. The first $d+1$ points define a unique weighted point y orthogonal to all p_{i_j} , $0 \leq j \leq d$. The test decides whether $\pi(p_{i_{d+1}}, y)$ is positive or negative. The case $\pi(p_{i_{d+1}}, y) = 0$ corresponds to the degeneracy when $p_{i_{d+1}}$ and y are orthogonal. As mentioned earlier this is assumed not to happen.

B.2 $\pi(p_{i_{d+1}}, y) < 0$ iff $\mathcal{M}_{0,1,\dots,d}^{i_0,i_1,\dots,i_d} \cdot \mathcal{M}_{0,1,\dots,d+1}^{i_0,i_1,\dots,i_{d+1}} < 0$.

The first minor is a corrective term that is necessary because $\pi(p_{i_{d+1}}, y)$ is unaffected if we change the order of the points although the sign of the second minor may change. We omit an argument for the correctness of B.2 as it will follow from the discussion of the more general attachment test below.

Lifted orthogonal point. Before discussing the size of a simplex σ_T and whether or not it is attached, we show how to determine the coordinates of the lifted version of the point $y = y_T$. Recall from section 3 that y is the point with minimum weight that is orthogonal to all points in T . As usual we write $y = (y', y'')$ and denote its coordinates as $y' = (v_1, v_2, \dots, v_d)$. Define $v_{d+1} = \sum_{j=1}^d v_j^2 - y''$, and $y^+ = (v_1, v_2, \dots, v_{d+1}) \in \mathbb{R}^{d+1}$, as in section 3. We show how to compute the coordinates of y^+ ; the location and weight of y can easily be derived from y^+ .

Let σ_T be a k -simplex, with $0 \leq k \leq d$, and let i_0, i_1, \dots, i_k be the indices of the points in T . The requirement that $\pi(p_\ell, y) = 0$, for $\ell = i_j$ and $0 \leq j \leq k$, can be formulated as follows.

$$\begin{aligned} 0 &= \pi(p_\ell, y) \\ &= |p'_\ell y'|^2 - p''_\ell - y'' \\ &= \sum_{m=1}^d (\phi_{\ell,m} - v_m)^2 - \left(\sum_{m=1}^d \phi_{\ell,m}^2 - \phi_{\ell,d+1} \right) - \left(\sum_{m=1}^d v_m^2 - v_{d+1} \right) \end{aligned}$$

$$= -2 \sum_{m=1}^d \phi_{\ell,m} v_m + \phi_{\ell,d+1} + v_{d+1}.$$

Thus, for $0 \leq j \leq k$, y^+ lies on the hyperplane $h_j : x_{d+1} = 2 \sum_{m=1}^d \phi_{\ell,m} x_m - \phi_{\ell,d+1}$, where $\ell = i_j$ as before. These $k+1$ hyperplanes intersect in a $(d-k)$ -flat in \mathbb{R}^{d+1} .

We need $d-k$ additional constraints to find the point y^+ on this $(d-k)$ -flat. These will be hyperplanes chosen so that each contains the points of $T' \in \mathbb{R}^d$ and is parallel to the $(d+1)$ st coordinate-axis. So their intersection is the $(k+1)$ -flat parallel to the $(d+1)$ st axis that contains the k -flat $\text{aff}(T')$. For each such hyperplane we arbitrarily pick $d-k-1$ coordinate-axes in addition to the $(d+1)$ st one. The requirement that the hyperplane be parallel to these $d-k$ axes determines it uniquely. Note that the number of choices is $\binom{d}{d-k-1} = \binom{d}{k+1} \geq d-k$, so it is indeed possible to find $d-k$ such hyperplanes. The constraint that a hyperplane be parallel to the j th coordinate-axis is equivalent to requiring it contains the point p_{-j} , which is at infinity in the direction of the j th axis. Let $I = \{i_{k+1}, i_{k+2}, \dots, i_{d-1}, i_d = -(d+1)\}$ be a set of $d-k$ indices between -1 and $-(d+1)$ inclusive; it represents a choice of $d-k$ axes including the $(d+1)$ st one. The corresponding hyperplane can be expressed as

$$\mathcal{M}_{0,1,\dots,d+1}^{i_0,\dots,i_k,i_{k+1},\dots,i_d,x} = 0. \quad (1)$$

Each row that corresponds to a negative index contains only one non-zero element, and this element is equal to 1. The absolute value of the minor is therefore unaffected if we delete these rows and the columns that correspond to the 1's. Let $J = \{j_0, j_1, \dots, j_{k+1}\} = \{0, 1, \dots, d+1\} - \{-i_{k+1}, -i_{k+2}, \dots, -i_d\}$ be the index set of the remaining columns, and notice that $0 \in J$ and $d+1 \notin J$. The equation

$$\mathcal{M}_{j_0,j_1,\dots,j_{k+1}}^{i_0,i_1,\dots,i_k,x} = \sum_{m=0}^{k+1} \mathcal{N}_{I,j_m} x_{j_m} = 0$$

defines the same hyperplane as (1), where

$$\mathcal{N}_{I,j_m} = (-1)^{k+1+m} \mathcal{M}_{j_0,j_1,\dots,\hat{j}_m,\dots,j_{k+1}}^{i_0,i_1,\dots,i_k}$$

for $0 \leq m \leq k+1$. The hat indicates that the marked index is dropped. Set $\mathcal{N}_{I,j} = 0$ for $j \notin J$. Putting things together we thus get $y^+ = (v_1, v_2, \dots, v_{d+1})$ as the unique solution to the linear system $\Gamma \cdot x = \gamma$:

$$\begin{pmatrix} 2\phi_{i_0,1} & 2\phi_{i_0,2} & \cdots & 2\phi_{i_0,d} & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 2\phi_{i_k,1} & 2\phi_{i_k,2} & \cdots & 2\phi_{i_k,d} & -1 \\ \mathcal{N}_{I_1,1} & \mathcal{N}_{I_1,2} & \cdots & \mathcal{N}_{I_1,d} & \mathcal{N}_{I_1,d+1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathcal{N}_{I_{d-k},1} & \mathcal{N}_{I_{d-k},2} & \cdots & \mathcal{N}_{I_{d-k},d} & \mathcal{N}_{I_{d-k},d+1} \end{pmatrix} \cdot \begin{pmatrix} x_1 \\ \vdots \\ x_{k+1} \\ x_{k+2} \\ \vdots \\ x_{d+1} \end{pmatrix} = \begin{pmatrix} \phi_{i_0,d+1} \\ \vdots \\ \phi_{i_k,d+1} \\ -\mathcal{N}_{I_1,0} \\ \vdots \\ -\mathcal{N}_{I_{d-k},0} \end{pmatrix}.$$

The sets I_1 through I_{d-k} are arbitrary but pairwise different choices of $d-k$ coordinate-axes including the $(d+1)$ st one, as described above. Let Γ_j be Γ after replacing the j th column with γ . By Cramer's rule we thus get $v_j = \frac{\det \Gamma_j}{\det \Gamma}$, for $1 \leq j \leq d+1$.

The size of a simplex. The size of a simplex σ_T, ϱ_T , is by definition the same as the weight of $y = y_T$, namely $\sum_{j=1}^d v_j^2 - v_{d+1}$. The following formula can now be readily derived.

$$\text{B.3 } \varrho_T = \frac{\sum_{j=1}^d (\det \Gamma_j)^2 - \det \Gamma \det \Gamma_{d+1}}{(\det \Gamma)^2}.$$

Attachment test. The k -simplex $\sigma_T \in \mathcal{R}$ is incident to various $(k+1)$ -simplices σ_U , and each U contains a point not in T . By 5.2, σ_T is attached iff $\pi(q, y_T) < 0$ for at least one of these points q . We thus define this test for an input of $k+2$ points with indices i_0, i_2, \dots, i_{k+1} between 1 and n inclusive. The first $k+1$ points span σ_T and thus define y . The question is whether or not $\pi(p_{i_{k+1}}, y) < 0$. We have

$$\begin{aligned} \pi(p_{i_{k+1}}, y) &= \sum_{j=1}^d (v_j - \phi_{i_{k+1},j})^2 - y'' - p_{i_{k+1}}'' \\ &= \sum_{j=1}^d (v_j - \phi_{i_{k+1},j})^2 - \left(\sum_{j=1}^d v_j^2 - v_{d+1} \right) - \left(\sum_{j=1}^d \phi_{i_{k+1},j}^2 - \phi_{i_{k+1},d+1} \right) \\ &= -2 \sum_{j=1}^d v_j \phi_{i_{k+1},j} + v_{d+1} + \phi_{i_{k+1},d+1} \\ &= \frac{1}{\det \Gamma} \left(-2 \sum_{j=1}^d \det \Gamma_j \phi_{i_{k+1},j} + \det \Gamma_{d+1} + \det \Gamma \phi_{i_{k+1},d+1} \right) \\ &= \frac{(-1)^{d+1}}{\det \Gamma} \cdot \det \begin{pmatrix} \phi_{i_0,d+1} & 2\phi_{i_0,1} & 2\phi_{i_0,2} & \dots & 2\phi_{i_0,d} & -1 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \phi_{i_k,d+1} & 2\phi_{i_k,1} & 2\phi_{i_k,2} & \dots & 2\phi_{i_k,d} & -1 \\ -\mathcal{N}_{I_1,0} & \mathcal{N}_{I_1,1} & \mathcal{N}_{I_1,2} & \dots & \mathcal{N}_{I_1,d} & \mathcal{N}_{I_1,d+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ -\mathcal{N}_{I_{d-k},0} & \mathcal{N}_{I_{d-k},1} & \mathcal{N}_{I_{d-k},2} & \dots & \mathcal{N}_{I_{d-k},d} & \mathcal{N}_{I_{d-k},d+1} \\ \phi_{i_{k+1},d+1} & 2\phi_{i_{k+1},1} & 2\phi_{i_{k+1},2} & \dots & 2\phi_{i_{k+1},d} & -1 \end{pmatrix}. \end{aligned}$$

Denote the final matrix by Λ . The sign does not change if we multiply the equation with the square of $\det \Gamma$, so we get

$$\text{B.4 } \pi(p_{i_{k+1}}, y) < 0 \text{ iff } (-1)^{d+1} \cdot \det \Gamma \cdot \det \Lambda < 0.$$

For the special case where $k = d$ we have $\det \Gamma = (-1)^{d+1} \cdot 2^d \cdot \mathcal{M}_{0,1,\dots,d}^{i_0,i_1,\dots,i_d}$ and $\det \Lambda = 2^d \cdot \mathcal{M}_{0,1,\dots,d+1}^{i_0,i_1,\dots,i_{d+1}}$. It follows that for this case B.2 and B.4 are equivalent.