ON THE SHAPE OF A SET OF POINTS IN THE PLANE[1]

by

H. Edelsbrunner[2], D.G. Kirkpatrick[3],
and R. Seidel[3]

Technical Report 81-15

September 1981

[1]Supported by the Natural Sciences and Engineering Research Council of Canada, grant A3583.

[2]Technical University Graz, Institut fuer Informationsverarbeitung, Steyrergasse 17, A-8010 Graz, Austria.

[3]University of British Columbia, Department of Computer Science, 2075 Wesbrook Mall, Vancouver, B.C., V6T 1W5, Canada.

ABSTRACT

A generalization of the convex hull of a finite set of points in the plane is introduced and analyzed. This generalization leads to a family of straight-line graphs, called "shapes", which seem to capture the intuitive notion of "fine shape" and "crude shape" of point sets.

Additionally, close relationships with Delaunay triangulations are revealed and, relying on these results, an optimal algorithm that constructs "shapes" is developed.

1. Introduction

The efficient construction of convex hulls for finite sets of points in the plane is one of the most exhaustively examined problems in the rather young field often referred to as "computational geometry". Part of the motivation is theoretical in nature. It seems to stem from the fact that the convex hull problem, like sorting, is easy to formulate and visualize. Furthermore, the convex hull problem, again like sorting, plays an important role as a component of a large number of more complex problems. Nevertheless, a lot of the work on convex hulls is motivated by direct applications in some of the more practical branches of Computer Science.

Akl and Toussaint [1], for instance, discuss the relevance of the convex hull problem to Pattern Recognition. By identifying and ordering the extreme points of a point set, the convex hull serves to characterize, at least in a rough way, the "shape" of such a set. Jarvis [9] presents several algorithms based on the so called nearest neighbour logic, that compute what he calls the "shape" of a finite set of points. The "shape", in Jarvis' terminology, is a notion made concrete by the algorithms that he proposes for its construction. Besides this lack of any analytic definition, the inefficiency of Jarvis' algorithms to construct the "shape" is a striking drawback. More recently, Fairfield [6] introduced a notion of shape of a finite point set based on the closest point Voronoi diagram of the set. He informally links his notion of shape with human perception but presents no concrete properties of his shapes, in particular no algorithmic results.

In this article, we introduce the notion of the "α-shape" of a finite set of points, for arbitrary real α. This notion is derived from a straightforward generalization of one common definition of the convex hull. Optimal algorithms for the construction of α-shapes and certain related structures are described. Consideration is given to the efficient construction of the α-shapes of a point set for several α's. The efficiency of our algorithms, in addition to other nice properties of α-shapes, leads us to believe that the family of all α-shapes, which we formalize as the shape spectrum, will find applications in a number of fields, including Pattern Recognition and Cluster Analysis.

In the next section, the notions of α-hull and α-shape are introduced along with a few of their basic properties. Section 3 describes the close connection between α-shapes and Delaunay triangulations. This serves as a basis for efficient algorithms to construct α-shapes and the shape spectrum presented in Section 4. In Section 5 we briefly discuss the problem of constructing an α-hull. The final section presents some concluding remarks and open questions.

2. Basic Notions

Given a set S of n points in the plane (n being a positive integer), we consider the following generalization of its convex hull. (The convex hull of S may be defined as the intersection of all closed halfplanes that contain all points of S.)

Definition 2.1:

Let α be an arbitrary, sufficiently small, positive real. The α-hull of S is the intersection of all closed discs with radius $1/\alpha$ which contain all points of S.

In order to achieve an intersection of discs, it has to be guaranteed that there exists at least one disc of the chosen size that contains all points. This implies that the smallest possible value for $1/\alpha$ is equal to the radius of the smallest enclosing circle. As a matter of fact, Jung [10] showed in 1901 that $1/\alpha$ no less than $3^{-1/2}$ times the diameter of S suffices, no matter how the points are distributed.

In Figure 2.1 the α-hull for some positive, sufficiently small α is depicted.

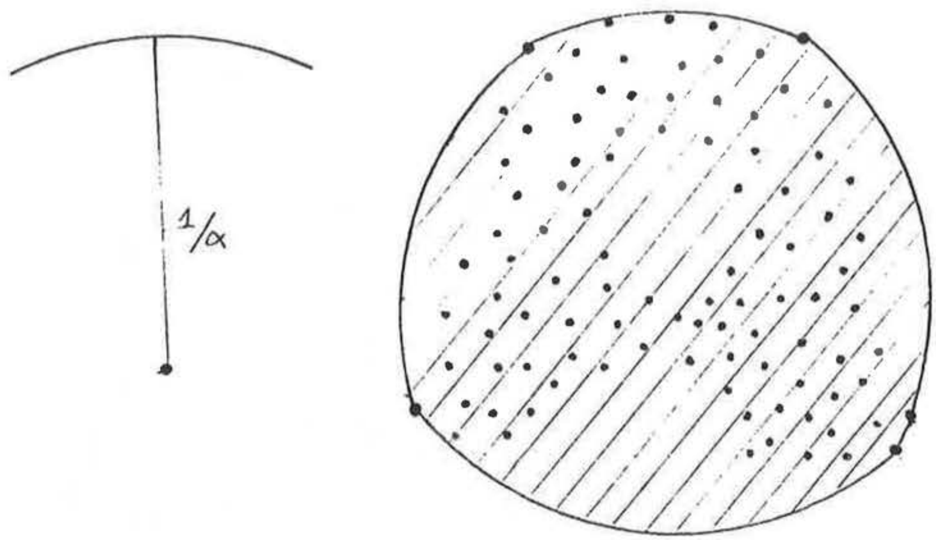

Figure 2.1: Positive α-hull.

Intuitively, large α (but still sufficiently small) give rise to hulls that have only in some sense "essential" extreme points on their boundary. For α going to zero, the α-hull approximates the common convex hull.

Definition 2.2:

For arbitrary negative reals α, the α-hull is defined as the intersection of all closed complements of discs (where these discs have radii $-1/\alpha$) which contain all points.

Figure 2.2 displays such a hull for the same point set as used in Figure 2.1.

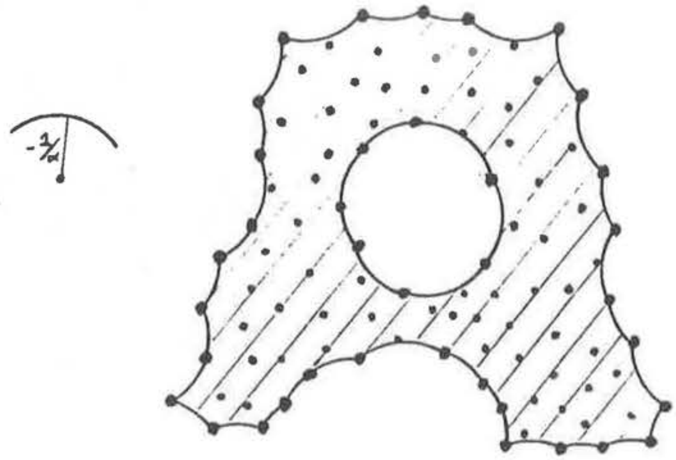

Figure 2.2: Negative α-hull.

For convenience let us define the <u>0-hull</u> as being the common convex hull of the points and let us agree that the intersection of <u>no</u> discs (which may occur for large positive α) is equal to the entire plane.

If we define a <u>generalized disc of radius $1/\alpha$</u> as a disc of radius $1/\alpha$ if $\alpha>0$, the complement of a disc of radius $-1/\alpha$ if $\alpha<0$, and a halfplane if $\alpha=0$, then the preceding definitions could be combined by saying: For an arbitrary real α and a set S of points in the plane, the <u>α-hull of S</u> is the intersection of all closed generalized discs of radius $1/\alpha$ which contain all points of S.

Thus we have a family of α-hulls for α ranging from $-\infty$ to ∞. Sample members of this family are the entire plane (for α sufficiently large), the smallest enclosing circle of S, the convex hull of S (for $\alpha=0$), and S itself (for α sufficiently small). All the members of this family satisfy the following simple relationship.

<u>Observation 2.1</u>:

The α_1-hull of a set of points is contained in the α_2-hull if $\alpha_1 \leq \alpha_2$.

Of central interest in this paper, however, will not be the continuous family of α-hulls, but the discrete family of what we call "α-shapes". Before defining α-shapes we need some additional notions.

<u>Definition 2.3</u>:

A point p in a set S is termed <u>α-extreme in S</u> if there exists a closed generalized disc of radius $1/\alpha$, such that p lies on its boundary and it contains all points of S. If for two extreme points p and q there exists a closed generalized disc of radius $1/\alpha$ with both points on its boundary and containing all other points, then p and q are said to be <u>α-neighbours</u>.

For convenience we assume that no four points in S are co-circular and no three points co-linear. The minor difficulties that arise in such cases can by treated by more elaborate definitions and considerations, which only tend to detract from our presentation.

Definition 2.4:

Given a set S of points in the plane and an arbitrary real α, the α-shape of S is the straight line graph whose vertices are the α-extreme points and whose edges connect the respective α-neighbours.

In Figure 2.3 the α-shape of the same set of points and the same reals α as used in Figure 2.1 and Figure 2.2 are displayed.

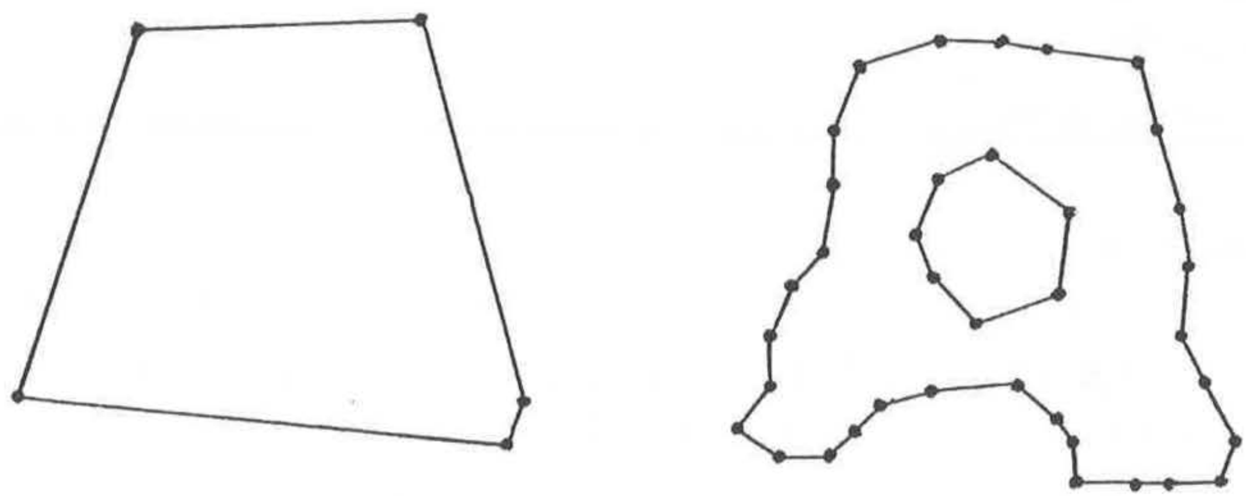

Figure 2.3: Positive and Negative α-shape.

The following corresponds directly with Observation 2.1:

Observation 2.2:

The set of α_1-extreme points in S is a subset of the α_2-extreme points if $\alpha_1 \geq \alpha_2$.

The α-hull was defined to be a straight line graph. However, in certain applications the intuitive notion of the "shape" of a set of points in the plane is not as well expressed by a set of straight line segments and points as by an area of "foreground" juxtaposed against a complementary "background". These two-dimensional notions can be captured by the α-shape by classifying some of its faces - it is a planar graph after all - as "interior" faces, that is "foreground", or as "exterior" faces, that is "background".

Definition 2.5:

Let S be a set of points in the plane and $\alpha \neq 0$. Let F be a face of the α-shape of S and e an edge incident to F.
For $\alpha > 0$ e is called a positive edge of F if the closed disc of radius $1/\alpha$, with the endpoints of e on its boundary, and its center strictly on the same side of e as F, contains all points of S. Otherwise e is called a negative edge of F.
For $\alpha < 0$ e is called a positive edge of F if the open disc of radius $1/\alpha$, with the endpoints of e on its boundary, and its center strictly on the same side of e as F, contains at least one point of S. Otherwise e is called a negative edge of F.

Definition 2.6:

For $\alpha \neq 0$ a face F of the α-shape of a planar point set S is called interior if one of its incident edges is a positive edge of F, and F is called exterior if one of its incident edges if a negative edge of F.

For $\alpha=0$ the bounded face of the 0-shape (i.e the convex hull) of S is the (only) interior face and the unbounded face is the (only) exterior face.

Figure 2.4 shows the α-shapes displayed in Figure 2.3 with their interior faces shaded. The observant reader will notice the similarity between interior faces of the α-shapes in Figure 2.4 and the α-hulls in Figure 2.1 and 2.2. In some sense the interior faces of an α-shape can be viewed as an α-hull with straight line segments as boundaries instead of circular arcs.

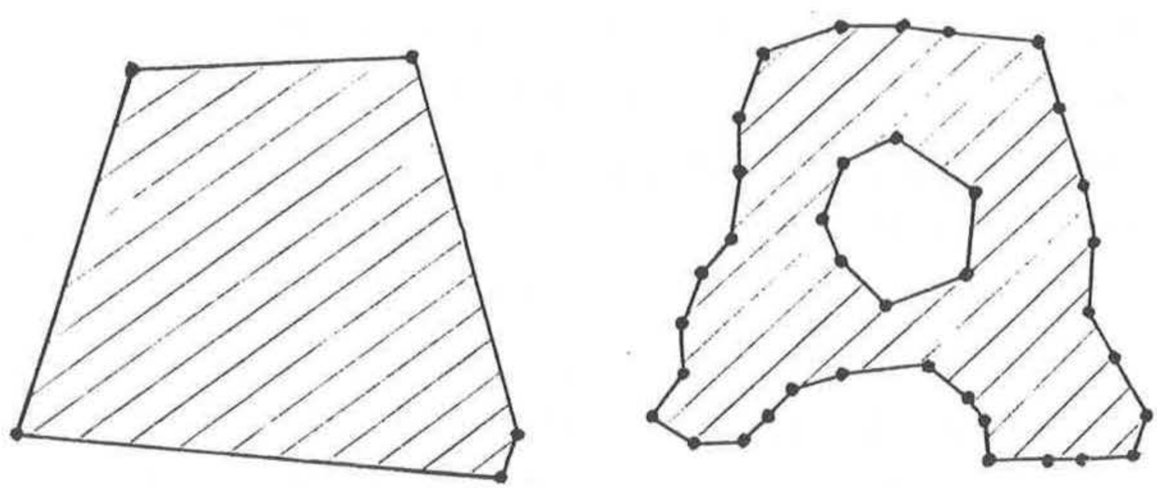

Figure 2.4: Interior Faces of α-shapes.

Intuitively, "relatively large" α tend to produce a rather crude shape of the points (the extreme being a chord or an inscribed triangle of the smallest enclosing circle), whereas smaller α reveal more and more details, until, as α approaches $-\infty$, all points are isolated extreme points of the shape. Thus α-neighbourliness is not monotonic with decreasing α like α-extremeness. However, as we shall see in the next section, two points can be α-neighbours for only some finite interval of α values. This, along with a characterization of exactly which pairs can be α-neighbours, is what makes possible the efficient construction of α-shapes.

3. α-Shapes and Delaunay Triangulations

In this section we make precise the rather close relationship that exists between α-shapes and Delaunay triangulations. Specifically we shall show that any α-shape of a set S of points is a subgraph of either the closest point or the furthest point Delaunay triangulation (whose definitions and properties are presented below). Other subgraphs of the closest point Delaunay triangulation have been studied, including the minimum spanning tree [16], the Gabriel graph [12], and the relative neighbourhood graph [18],[17]. However in general none of these graphs is a member of the family of α-shapes of S.

First we give a few facts about Voronoi diagrams and Delaunay triangulations. Given a set S of n points in the plane the closest point Voronoi diagram of S $VD_c(S)$ is a partition of the plane into n regions $V_p, p \in S$, where

$$V_p = \{x \mid d(p,x) < d(q,x),\ p \neq q \in S\} .$$

Similarly the furthest point Voronoi diagram of S $VD_f(S)$ is defined by the regions

$$W_p = \{x \mid d(p,x) > d(q,x),\ p \neq q \in S\},\ p \in S.$$

We will need the following properties of these diagrams:

Fact 3.1:

The regions V_p and W_p are convex and bounded by straight line segments, called Voronoi edges, for all $p \in S$.

Fact 3.2:

Each region V_p of $VD_c(S)$ contains p. Provided n>1, each region W_p of $VD_f(S)$ does not contain p.

Fact 3.3:

The regions V_p and W_p are unbounded iff p is a convex hull point of S. Otherwise V_p is a non empty convex polygon and W_p is empty.

Two points p and q of S are said to be closest (respectively furthest) point Voronoi neighbours if the boundaries of V_p and V_q (resp. W_p and W_q) have a common point.

Fact 3.4:

Two points p and q of S are closest and furthest point Voronoi neighbours iff (p,q) is a convex hull edge of S.

The closest (resp. furthest) point Delaunay triangulation of S, $DT_c(S)$ (resp. $DT_f(S)$), is defined as the straight line dual of $VD_c(S)$ (resp. $VD_f(S)$); i.e. there is a straight line edge between p and q iff they are closest (resp. furthest) point Voronoi neighbours.

Fact 3.5:

Both the closest and furthest point Voronoi diagram (as well as the respective Delaunay triangulations) of n points can be constructed in O(nlogn) time and O(n) space. Furthermore the closest or furthest point Voronoi diagram can be constructed from the respective Delaunay triangulation in O(n) time, and vice versa.

For proofs leading to facts 3.1 to 3.5 and other properties of these constructs consult [15]. An algorithm which unifies the closest and furthest point case is given by Brown [2].

In the following we assume that our point set S is fixed. The relationship between the Delaunay triangulations and α-shapes is given by the following lemma.

Lemma 3.1:

The α-shape of S is a subgraph of $DT_f(S)$ if $\alpha \geq 0$ and a subgraph of $DT_c(S)$ if $\alpha \leq 0$.

Proof:

Trivially each vertex of an α-shape is also a vertex of the respective Delaunay triangulation. Next, we need to show that, if p and q are α-neighbours, then they are adjacent in the respective Delaunay triangulation. We consider three cases:

(i) $\alpha=0$: The convex hull is a subgraph of both $DT_c(S)$ and $DT_f(S)$ by Fact 3.4.

(ii) $\alpha>0$: Let p and q be α-neighbours and let c be the center of the disc of radius $1/\alpha$ which touches p and q and contains all other points $x \in S$. Clearly $d(c,p) \geq d(c,x)$ and $d(c,q) \geq d(c,x)$ for all $x \in S$, $p \neq x \neq q$. As $d(c,p) = d(c,q)$ both p and q are furthest neighbours of c. Hence p and q must be furthest point Voronoi neighbours and therefore p and q must be adjacent in $DT_f(S)$.

(iii) $\alpha<0$: The proof is essentially the same as in (ii) replacing "furthest" by "closest", "contains all" by "contains no", "$\geq$" by "$\leq$", and "$DT_f(S)$" by "$DT_c(S)$".

Q.E.D.

The following two lemmas are important for the construction of an α-shape. They tell for which $\alpha \in R$ a vertex or an edge of a Delaunay triangulation is also a vertex or edge of the α-shape.

<u>Lemma 3.2</u>:

For each point $p \in S$ there exists a real number $\alpha\text{-max}(p)$ such that p is α-extreme in S if and only if $\alpha \leq \alpha\text{-max}(p)$.

Proof:

For the proof of the lemma we have two cases to consider:

(i) p lies on the convex hull of S:

Recall the definition of α-extremeness for positive α: p must lie on the boundary of a disc of radius $1/\alpha$ containing all remaining points of S. The center of such a disc has to lie in the furthest point Voronoi region W_p of p. It is not difficult to see that W_p actually comprises exactly all possible centers of discs touching p and containing S. W_p is an unbounded convex region which does not contain p (if one disregards the trivial case of $|S|=1$). Therefore there are discs of radius r touching p and containing S exactly for all $r \geq d(W_p,p)$. Thus p is α-extreme for $0 < \alpha \leq 1/d(W_p,p)$. A convex hull point is trivially α-extreme for non-positive α, hence p is α-extreme for all $\alpha \leq 1/d(W_p,p) =: \alpha\text{-max}(p)$.

(ii) p is not a convex hull point of S:

It is easy to see that p cannot be α-extreme for $\alpha \geq 0$. For p to be α-extreme for negative α, p has to lie on the boundary of a disc of radius $-1/\alpha$ containing none of the remaining points of S. The set of centers of such discs is exactly the closest point Voronoi region V_p of p. By Facts 3.2 and 3.3 V_p is a convex polygon containing p. Therefore there are discs of radius r touching p and not containing S exactly for all $r \leq d_p := \max\{d(p,x) \mid x \in V_p\}$. This implies that p is α-extreme for all $\alpha \leq -1/d_p =: \alpha\text{-max}(p)$.

Q.E.D.

Note that if $DT_c(S)$ and $DT_f(S)$ are given, the set $\{\alpha\text{-max}(p) | p \in s\}$ can be computed in linear time by testing for each point of S the distances to the centers of the circumscribed circles of the incident triangles.

Lemma 3.3:

For every edge e belonging to either $DT_c(S)$ or $DT_f(S)$ there are real numbers α-min(e) and α-max(e), $\alpha\text{-min}(e) \leq \alpha\text{-max}(e)$, such that e is an edge of the α-shape of S if and only if $\alpha\text{-min}(e) \leq \alpha \leq \alpha\text{-max}(e)$.

Proof:

First we state without proof the following two facts:

a) Given a point p and a semi-infinite line segment s there exists a positive real number a=a(p,s) such that

$\{d(p,x) | x \in s\} = [a, \infty)$.

b) Given a point p and a closed line segment s there positive real number a=a(p,s) and b=b(p,s) such that

$\{d(p,x) | x \in s\} = [a,b]$.

Now, let p and q be the two points incident to an edge e. We have to consider three cases:

(i) e is a non convex hull edge of $DT_c(S)$:

The center of a disc touching p and q and not containing other points of S must lie on the bisector between p and q, and must be closer to p and q than to any other point of S. The locus of points having exactly this property is the straight line segment separating the Voronoi regions V_p and V_q (i.e. the Voronoi edge dual to e). Thus by Fact b) there are discs of radius r touching p and q and not containing other points of S exactly for all r, $a(p,s) \leq r \leq b(p,s)$. It is easy to see that as a consequence of Fact 3.4 there are no discs touching p and q and containing the remaining points of S. Thus e is an edge of the α-shape exactly for all α, $-1/a(p,s) =: \alpha\text{-min}(e) \leq \alpha \leq \alpha\text{-max}(e) := -1/b(p,s)$.

(ii) e is a non convex hull edge of $DT_f(S)$:

The proof is essentially the same as in (i) replacing "furthest" by "closest", "contains all" by "contains no", etc.

(iii) e is a convex hull edge:

First note that p and q are trivially α-neighbours for α=0. The locus of all centers of discs touching p and q and containing all other points of S is exactly the closed semi-infinite line segment w separating the furthest point Voronoi regions W_p and W_q. Thus by fact b) there are discs of radius r touching p and q and containing all other points of S exactly for all $r \geq a(p,w)$.

By the same argument there are discs of radius r touching p and q and containing none of the other points of S exactly for $r \geq a(p,v)$, where v is the semi-infinite line segment separating V_p and V_q.

Thus p and q are α-neighbours, i.e. e is an edge of the α-shape, for all α,

$$-1/a(p,v) =: \alpha\text{-min}(e) \leq \alpha \leq \alpha\text{-max}(e) := 1/a(p,w).$$

Q.E.D.

Again, note that if $DT_c(S)$ and $DT_f(S)$ are given, α-min(e) and α-max(e) can be computed in constant time for each edge e considering the centers of the circumscribed circles of the two incident triangles only.

4. Construction of α-Shapes and the Shape Spectrum

4.1 α-Shapes

Together, lemmas 3.1, 3.2, and 3.3 give rise to the following algorithm to determine the α-shape of a set S.

Algorithm 4.1 (Construction of the α-shape of S)

(1) Construct a triangulation DT:

if $\alpha \geq 0$, construct $DT_f(S)$

if $\alpha < 0$, construct $DT_c(S)$

(2) Determine the α-extreme points of S

The information provided by DT suffices for this task, see also Lemma 3.1 and 3.2.

(3) Determine the α-neighbours of S

Again, DT contains all the information necessary to perform this task, see also Lemma 3.1 and 3.3.

(4) Output the α-shape

Output the graph on the α-extreme points with all α-neighbour connections.

The correctness of Algorithm 4.1 follows immediately from Lemmas 3.1, 3.2, and 3.3. A straightforward analysis of Algorithm 4.1 gives rise to the following:

Theorem 4.1:

The α-shape of n points in the plane can be determined for an arbitrary real α in time O(nlogn) and space O(n).

Proof:

It suffices to show that the stated bounds hold for Algorithm 4.1.

Step (1) can be done in O(nlogn) time and O(n) space by Fact 3.5. Once the appropriate Delaunay triangulation has been constructed steps (2),(3), and (4) can be done (see the notes following Lemmas 3.2 and 3.3) in O(n) time and O(n) space. Whenever in step (2) or (3) the actual value of α-max(p), or α-min(e), or α-max(e) cannot be computed (because p is a convex hull point, or e is a convex hull edge) the value 0 can be used as an appropriate substitute.

Q.E.D.

4.2 Interior and Exterior Faces

It should be clear that Algorithm 4.1 can be generalized quite easily to yield, in addition to the α-shape, its interior and exterior faces. However, a few remarks about the properties of interior and exterior faces seem to be appropriate at this point. Their rather straightforward but lengthy proofs are omitted.

1. A face F of an α-shape of a point set S is either an interior face, that is all its bounding edges are positive edges of F, or it is an exterior face, that is all its incident edges are negative edges of F. The only minor exception (i.e. faces that are both interior and exterior) are faces which are triangles with circumscribed circle of radius exactly $1/|\alpha|$ and with center outside the triangle. This situation reflects a non-continuous change in the α-hull for varying α at such values.

2. Interior faces do not properly contain α-extreme points.

3. For negative α, any closed disc of radius $-1/\alpha$ with center in an interior face of an α-shape of a set S contains a point of S. This means that interior faces represent clusters of S.

4. For $\alpha \geq 0$ (resp. $\alpha \leq 0$) the interior faces of the α-shape of S are exactly the union of the triangles in $DT_f(S)$ (resp. $DT_c(S)$) whose circumscribed circles have radius not greater than $1/|\alpha|$. Thus the interior faces of an α-shape can be trivially computed from the appropriate Delaunay triangulation in linear time without constructing the α-shape itself.

4.3 The Shape Spectrum

It is easy to envision applications in which the α-shape of a point set is desired for a number of different α's. As the analysis of Algorithm 4.1 makes clear, it is possible to construct α-shapes, following an initial expenditure of O(nlogn) to contruct both Delaunay triangulaions, at a cost of O(n) per shape. In fact, as we shall see, a slightly tighter bound is possible by a careful choice of data structures. As an intermediate step in this construction, and because it is an interesting entity in its own right, we consider first what we call the shape spectrum of a point set S.

Definition 4.1: The shape spectrum SP(S) of a point set S is defined to be the set of intervals int(p) = (-∞ ,α-max(p)] and int(e) = [α-min(e),α-max(e)], p ε S, and e an edge of $DT_f(S)$ or $DT_c(S)$ of S.

The shape spectrum of a point set can be seen as an encoding of all possible α-shapes of that set. As the following lemma shows it also has the nice property that it is no more difficult to construct than the α-shape for a single fixed α.

Lemma 4.1:

The shape spectrum SP(S) of a set S of n points can be constructed in time O(nlogn) and space O(n).

Proof: Immediate generalization of Algorithm 4.1. Q.E.D.

Given the spectrum SP(S) of a set S, a number of problems concerning α-shapes of S can be solved with surprising efficiency.

1. The most prominent, of course, is, given SP(S) and some α_o, find the α_o-shape of S. This trivially can be done in linear time by determining all points p and edges e such that $\alpha_o \in int(p)$ and $\alpha_o \in int(e)$. However, by using a more advanced data structure to store the intervals of SP(S), such as Edelsbrunner's interval tree [5], called tile tree in the independent paper of McCreight [13], the α_o-shape of S can actually be constructed in time O(logn + t), where t is the number of points and edges in the α_o-shape.

2. It may be useful in certain applications to find an α-shape satisfying certain properties. For example, suppose, one wants to find an α_o such that the α_o-shape of S contains exactly k points. If the endpoints of the intervals int(p), $p \in S$, are stored in a sorted array, α_o can clearly be found in O(logn) time.

3. A similar problem addresses the fine tuning of α-shapes: given the α_o-shape of S for some α_o, find the largest $\alpha_1 < \alpha_o$ (or the smallest $\alpha_1 > \alpha_o$), such that the α_1-shape is different from the α_o-shape of S. By maintaining a pointer into the sorted list of the endpoints of the intervals in SP(S), this question can be answered in constant time.

4. An inverse problem to the construction of α-shapes asks for a given graph G on a subset of S, whether G is an α-shape of S for some α. The answer to this question can be found in linear time by the following procedure which uses a sorted list L of the endpoints of the intervals in SP(S). First confirm that each edge e of G is a Delaunay edge, that is, int(e) is defined. Initialize three counters i,j, and k to zero and scan L in decreasing order. If at any point during this scan i equals the number of vertices in G, j equals the number of edges of G, and k equals zero, then G is an α-shape of S. If an element of L being scanned is the right endpoint of an interval int(x) ε SP(S), increment i if x is a vertex of G, increment j if x is an edge in G, and increment k otherwise. If an element of L being scanned is the left endpoint of an interval int(x) ε SP(S), decrement k if x is an edge not in G, and stop otherwise, because in this case G cannot be an α-shape of S.

5. Constructing the α-Hull

In the preceding section we presented an O(nlogn) algorithm for the construction of an α-shape of a set S of n points in the plane. We went on to define the spectrum SP(S). As SP(S) contains only linearly many elements we can argue that for a given set S there are at most linearly many distinct α-shapes. If we turn our attention to α-hulls the situation becomes quite different. The number of distinct α-hulls is uncountable because for every two distinct $\alpha_1, \alpha_2 \in (-1/a, 1/b]$, where a is the radius of the smallest circle defined by three points in S, and b is the radius of the smallest enclosing circle of S, the α_1-hull is different from the α_2-hull of S. So it is quite surprising that α-hulls can be constructed efficiently for any real α. Specifically we shall show that for any real α the α-hull of S has a linear description and can be constructed in O(nlogn) time. As these facts seem to be quite obvious for $\alpha \geq 0$, we will concern ourselves only with the case of negative α.

At first let us recall the definition of the α-hull of S for negative α: it is defined as the intersection of all complements of open discs of radius $-1/\alpha$ which contain no point of S. By DeMorgan's law an equivalent definition is that the α-hull is the complement of the union of all open discs of radius $-1/\alpha$ which contain no point of S. Because of the fact that a disc of radius R can be represented as the union of open discs of radius $r \leq R$, there is another equivalent definition for the α-hull which we find more convenient to work with:

> The α-hull ($\alpha < 0$) of S is the complement of the union of all open discs of radius not less than $-1/\alpha$ which contain no point of S.

Our main problem now is to determine the union of all these discs. The set of all open discs of radius not less than $-1/\alpha$ is still rather unwieldy, but fortunately, as the next lemma shows, we can restrict our attention to a much smaller set of open discs.

In the following let $S(x,r)$ denote the open disc of radius r centered at x.

Lemma 5.1:

Let D be an open disc which does not contain any points of S. Either D lies entirely outside the convex hull of S or there is an open disc D_1 which contains D but no points of S and which has its center on an edge of $VD_c(S)$.

Proof:

Let $D = S(c,r)$ be a disc which does not contain any points of S. Let $p \in S$ be the point, such that $d(c,p) = \min\{d(c,x) \mid x \in S\}$. Clearly the disc $D' = S(c,d(c,p))$ touches p but does not contain any point of S. Furthermore $D \subset D'$. Let h be the straight line through c and p and let t be the intersection of h with the bounding edge of V_p such that c lies on the closed line segment between p and t. (If such a t does not exist, D' and D lie entirely outside the convex hull of S.) Clearly the open disc $D_1 = S(t,d(t,p))$ has the desired properties; i.e. D_1 contains no point of S and has its center on an edge of $VD_c(S)$, and $D \subset D' \subset D_1$.

Q.E.D.

As a consequence of Lemma 5.1 the α-hull ($\alpha<0$) of S can be expressed as the complement of the union of open discs of radius not less than $-1/\alpha$ which do not contain any points of S and which have their centers of the edges of $VD_c(S)$. Next we state without proof two basic geometric facts which will allow us to consider an even smaller set of discs.

Fact 5.1:

Let p and q be two distinct points in the plane and let L be a closed line segment on the bisector between p and q which is bounded by the points a and b. Then

$$\bigcup \{S(x,d(x,p)) \mid x \in L\} = S(a,d(a,p)) \cup S(b,d(b,p)).$$

Fact 5.2:

Let p and q be two distinct points in the plane and let L be a semi-infinite closed line segment on the bisector between p and q which is bounded by point a. Then

$$\bigcup \{S(x,d(x,p)) \mid x \in L\} = S(a,d(a,p)) \cup H(p,q),$$

where H(p,q) denotes the open halfplane defined by the straight line through p and q which contains the infinite portion of L.

Lemma 5.2:

For negative α the α-hull of a set S of n points can be expressed as the complement of the union of O(n) open discs and halfplanes.

Proof:

As we remarked after Lemma 5.1 we only have to consider appropriate discs centered on edges of $VD_c(S)$. Let $p,q \in S$ be two Voronoi neighbours and let r be the edge separating V_p and V_q. Clearly for every $x \in r$, S(x,d(x,p)) contains no points of S. Now observe that the set $\{x \in r \mid d(x,p) \geq -1/\alpha\}$ is either empty or forms one or two closed line segments. (If p and q are convex hull points, than one line segment is semi-infinite.) Thus by Facts 5.1 and 5.2 the union of open discs centered on r and of radius not less than $-1/\alpha$ and which contain no points of S can be expressed as the union of at most 4 open discs or halfplanes. As the number of edges in $VD_c(S)$ is linear in n, one can conclude that for negative α the complement of the α-hull of S can be expressed as the union of O(n) open discs and halfplanes.

Q.E.D.

With this result we can easily prove the following:

<u>Theorem 5.1</u>:

The α-hull of a set S of n points can be computed in time O(nlogn) using O(n) space.

Proof:

We consider two cases:

(i) $\alpha<0$: By Lemma 5.2 it suffices to find the union of O(n) discs and halfplanes which can be determined from $VD_c(S)$ in O(n) time. It is not difficult to see that the union of the halfplanes in question is the complement of the convex hull of S which can be determined from $VD_c(S)$ in linear time. The union of the O(n) discs can be constructed O(nlogn) time using a method developed by Brown [3]. Special care must be taken to identify isolated points of the α-hull. This can be done in a way similar to step (2) of Algorithm 4.1.

(ii) $\alpha \geq 0$: For non-negative α the α-hull can be derived directly and in linear time from the α-shape. In certain applications straightforward generalizations of common convex hull algorithms, as for instance Graham's [7], may be preferred. The details are left to the reader.

Q.E.D.

6. Conclusions

In this paper we developed the notion of the α-hull and α-shape of a set of points in the plane. We presented efficient algorithms to construct α-shapes and α-hulls which are based on the intimate relation of these constructs with Delaunay triangulations. We introduced the notion of the shape spectrum and briefly discussed some of its applications. Because α-shapes have nice theoretical properties and can be constructed efficiently, and because of the fact that they seem to capture the intuitive notion of "finer" or "cruder shape" of a planar pointset, we contend that α-shapes will be of good use in practical applications.

In conclusion we want to discuss a few related problems and point out some generalizations.

At first we briefly address the question of dynamization: given the α-shape of a set S for some α, how does the insertion of a point into S or the deletion of a point from S affect the α-shape? As Voronoi diagrams can be updated in linear time [7],[14], and α-shapes can be constructed from the Voronoi diagrams in linear time, the update time for α-shapes is O(n). This is even true for the shape spectrum, as long as it is just treated as a set of intervals. But we have not been able to design a linear time update algorithm which also maintains any of the additional data structures on SP(S) (such as the interval tree or the sorted lists) which we mentioned in Section 4.

Next we want to point out that the notion of α-shapes generalizes nicely to point sets in 3-space or in k-space. One can define α-extreme points, α-neighbours and α-triples, etc., similarly as in the definitions of Section 2, by using balls of radius 1/α instead of discs. The 3D-α-shape is related to 3D-Voronoi diagrams in a similar way as 2D-α-shapes to 2D-Voronoi diagrams. Lemmas 3.1 to 3.3 and Algorithm 4.1 carry over to 3D without much modification. But because Voronoi diagrams in 3-space can be very complex and can incorporate $O(n^2)$ Voronoi neighbours ([15],[11], [4]), one cannot expect to find an algorithm which uses less than quadratic time in the worst case.

REFERENCES

[1] Akl, S.G. and Touissant, G.T.
Efficient Convex Hull Algorithms for Pattern Recognition Applications.
Proc. 4th Intern. Joint. Conf. on Pattern Recognition, Kyoto, pp.1-5, 1978.

[2] Brown, K.Q.
Voronoi Diagrams from Convex Hulls.
Info. Proc. Letters 9, pp.223-228, 1979.

[3] Brown, K.Q.
Geometric Transforms for Fast Geometric Algorithms.
Carnegie-Mellon Univ., Pittsburgh, Pennsylvania, Dept. of Comp. Sci., Rep. CMU-CS-80-101, 1980.

[4] Dewdney, A.K. and Vranch, J.K.
A Convex Partition of R^3 with Applications to Crum's Problem and Knuth's Post Office Problem.
Utilitas Math. 12, pp.193-199, 1977.

[5] Edelsbrunner, H.
A New Approach to Rectangle Intersections.
Submitted to Computing.

[6] Fairfield, J.
Contoured Shape Generation: Forms that People see Dot Patterns.
Proc. IEEE Conf. on Cybernetics and Society 1979, pp.60-64.

[7] Gowda, I.G. and Kirkpatrick, D.G
Exploiting Linear Merging and Extra Storage in the Maintainance fo Fully Dynamic Geometric Data Structures.
Proc. 18th annual Allerton Conf on Communication, Control, and Computing, pp.1-10, October 1980.

[8] Graham, R.L.
An Efficient Algorithm for Determining the Convex Hull of a Finite Planar Set.
Info. Proc. Letters 1, pp.132-133, 1972.

[9] Jarvis, R.A.
Computing the Shape Hull of Points in the Plane.
Proc. of the IEEE Comp. Soc. Conf. on Pattern Recognition and Image Processing, pp.231-241, 1977.

[10] Jung, H.W.E.
Ueber die kleinste Kugel, die eine raeumlich Figur einschliesst.
Journal f.d. Reine u. Angew. Math. 123, pp.241-257, 1901.

[11] Klee, V.E.
On the Complexity of d-dimensional Voronoi Diagrams.
Archiv d. Math., vol. 34, pp.75-80, 1980.

[12] Matula, D.W. and Sokal, R.R.
Properties of Gabriel Graphs Relevant to Geographic Variation Research and the Clustering of Points in the Plane.
Geog. Analysis, vol. 12, pp.205-222, July 1980.

[13] McCreight, E.M.
Efficient Algorithms for Enumerating Intersecting Intervals and Rectangles.
Palo Alto Res. Center, Palo Alto, California, Rep. CSL-80-9, 1980.

[14] Overmars, M.H.
Dynamization of Order Decomposable Set Problems.
Univ. of Utrecht, Dept. of Comp. Sci., Tech. Rep. RUU-CS-80-9, 1980.

[15] Preparata, F.P.
Steps into Computational Geometry.
Univ. of Illinois, Urbana, Illinois, Coordinated Sci. Lab., Rep. R-792, 1977.

[16] Shamos, M.I.
Geometrical Complexity.
Proc. 7th ACM Symp. on Theory of Computing, pp.224-233, 1975.

[17] Supovit, K.J.
The Relative Neighbourhood Graph with an Application to Minimum Spanning Trees.
Tech. Rep., Dept. of Comp. Sci., Univ. of Illinois, Urbana, August 1980.

[18] Touissant, G.T.
The Relative Neighbourhood Graph of a Finite Planar Set.
McGill Univ., Montreal, Tech. Rep. SOCS 79.9, May 1979.