ACKNOWLEDGMENT

The authors express their appreciation for numerous constructive suggestions, which led to improvements on various phases of the manuscript, to Dr. Marvin Simon of JPL and to Professor George L. Turin of University of California, Berkeley. The junior author also gratefully acknowledges the role of the latter as her M.S. research advisor on the project which formed the nucleus of this work.

REFERENCES

[1] I. Gurantz, S. Blake, E. Hoversten, and J. Petranovich, "A high performance multiple data rate burst modem for satellite packet communication," EASCON Conference Record, Nov. 1981.
[2] C. Heegard, J. A. Heller, and A. J. Viterbi, "A microprocessor-based PSK modem for packet transmission over satellite channels," *IEEE Trans. Commun.*, vol. COM-26 no. 5 pp. 552–564, May, 1978.
[3] A. J. Viterbi, *Principles of Coherent Communication*. New York: McGraw-Hill, 1966.
[4] G. C. Clark, Jr., and J. B. Cain, *Error-Correction Coding for Digital Communication*. New York: Plenum Press, 1981.
[5] A. J. Viterbi and J. K. Omura, *Principle of Digital Communication and Coding*. New York: McGraw-Hill, 1979.
[6] J. J. Stiffler, *Theory of Synchronous Detection*. Englewood Cliffs, NJ: Prentice-Hall, 1971, p. 247.
[7] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products*. New York: Academic, 1965.
[8] M. Abramowitz and I. A. Stegun, Eds., *Handbook of Mathematical Functions*. Washington, DC: National Bureau of Standards, 1964.

On the Shape of a Set of Points in the Plane

HERBERT EDELSBRUNNER, DAVID G. KIRKPATRICK, AND RAIMUND SEIDEL

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, "α-shapes," which seem to capture the intuitive notions of "fine shape" and "crude shape" of point sets. It is shown that α-shapes are subgraphs of the closest point or furthest point Delaunay triangulation. Relying on this result an optimal $O(n \log n)$ algorithm that constructs α-shapes is developed.

I. 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, much of the work on convex hulls is motivated by direct applications in some of the more practical branches of computer science.

Manuscript received Sept. 9, 1981; revised June 1982. Preliminary results from this work were presented at the Conference on Graph Theory in Linz, June 1981. The work was supported by the National Sciences and Engineering Research Council of Canada, Grant A3583.

H. Edelsbrunner is with the Technical University of Graz, Institute fuer Informationsverarbeitung, Schiesstattgasse 4a, A-8010 Graz, Austria.

D. G. Kirkpatrick and R. Seidel are with the University of British Columbia, Department of Computer Science, 2075 Wesbrook Mall, Vancouver, BC, V6T 1W5 Canada.

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 neighbor 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 the 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. (See also Toussiant [20] for another definition of the shape of a set based on the Voronoi diagram.)

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 III 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 IV. In Section V we briefly discuss the problem of constructing an α-hull. The final section presents some concluding remarks and open questions.

II. Basic Notions

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

Definition 1: Let α be a sufficiently small but otherwise arbitrary *positive* real. The *α-hull* of S is the intersection of all closed discs with radius $1/\alpha$ that contain all the 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. In 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 Fig. 1 the α-hull for a particular sufficiently small positive α is depicted. Intuitively, large (but still sufficiently small) α give rise to hulls that have only in some sense "essential" extreme points on their boundary. As α approaches zero, the α-hull approximates the common convex hull.

Definition 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 $-1/\alpha$) that contain all the points of S.

Fig. 2 displays such a hull for the same point set as used in Fig. 1. For convenience, let us define the *0-hull* as being the usual convex hull of the points and let us agree that the intersection of *no* discs (which may occur for large positive α) is equal to the entire plane.

If we define a *generalized disc of radius* $1/\alpha$ 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 as follows: for an arbitrary real α and a set S of points in the plane, the *α-hull of* S is the intersection of all closed generalized discs of radius $1/\alpha$ that contain all the 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 (when $1/\alpha$ equals its radius), 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.

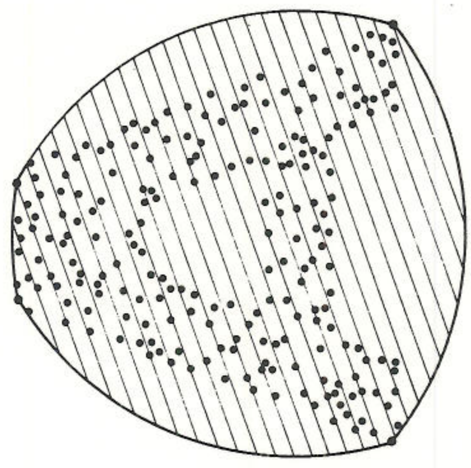

Fig. 1. Positive α-hull.

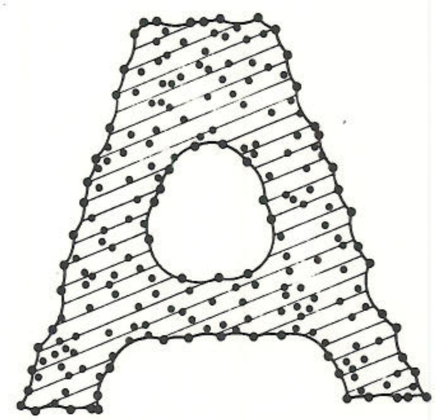

Fig. 2. Negative α-hull.

Observation 1: The α_1-hull of a set of points is contained in the α_2-hull if $\alpha_1 \leqslant \alpha_2$. (For example, the α-hull of Fig. 2 is contained in that of Fig. 1.)

The focus of this paper, however, will not be the continuous family of α-hulls, but rather the discrete family of what we call "α-shapes."

Before defining α-shapes we need some additional notions.

Definition 3: A point p in a set S is termed *α-extreme in* S if there exists a closed generalized disc of radius $1/\alpha$, such that p lies on its boundary and it contains all the points of S. If for two α-extreme points p and q there exists a closed generalized disc of radius $1/\alpha$ with both points on its boundary and which contains all other points, then p and q are said to be *α-neighbors*.

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

Definition 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-neighbors.*

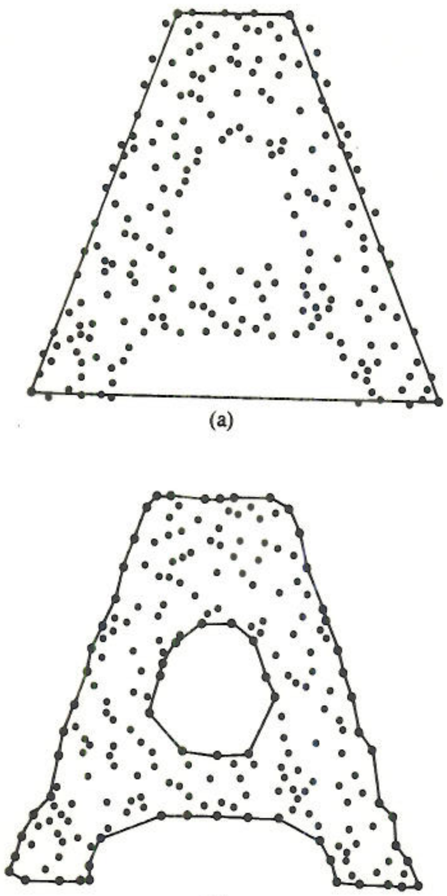

(a)

(b)

Fig. 3. Positive and negative α-shape.

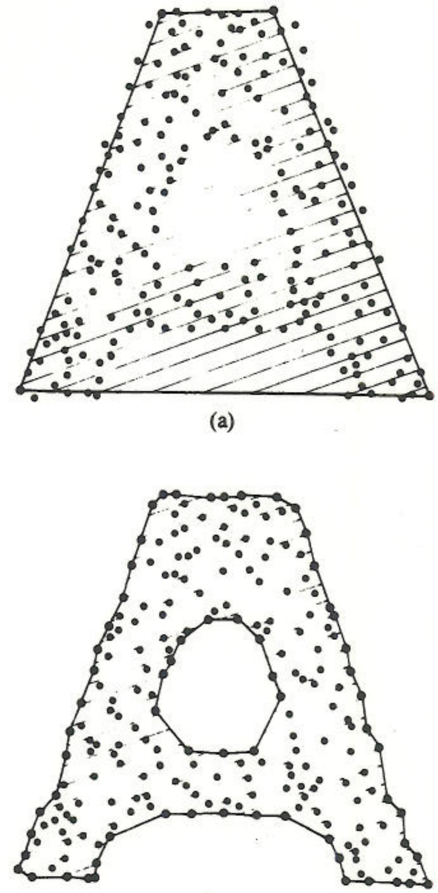

(a)

(b)

Fig. 4. Interior faces of α-shapes.

In Fig. 3 the α-shapes of the same set of points and the same reals α as used in Fig. 1 and Fig. 2 are displayed.

The following observation corresponds directly with Observation 1.

Observation 2: The set of α_1-extreme points in S is a subset of the α_2-extreme points if $\alpha_1 \geqslant \alpha_2$. For example, the vertices of the α-shape in Fig. 3(a) are a subset of those in Fig. 3(b).

The α-shape was defined to be a straight line graph. In certain applications, however, 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 with the α-shape by classifying some of its faces—it is a planar graph after all—as either "interior" (foreground) or "exterior" (background) faces.

Definition 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 let e be a boundary edge of 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 the 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 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 boundary edges is a positive edge of F, and F is called *exterior* if one of its boundary edges is 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.

Fig. 4 shows the α-shapes displayed in Fig. 3 with their interior faces shaded. Note the similarity between the interior faces of the α-shapes in Fig. 4 and the α-hulls in Figs. 1 and 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.

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-neighborliness is not monotonic with decreasing α like α-extremeness. As we shall see in the next section, two points can be α-neighbors only for some finite interval of α values. This, along with a characterization of exactly which pairs can be α-neighbors, is what makes possible the efficient construction of α-shapes.

III. 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 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 [17], the Gabriel graph [12], and the relative neighborhood graph [19], [18]. In the general case, however, none of these graphs is a member of the family of α-shapes of S.

First we present 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 *covering* of the plane by n regions V_p, $p \in S$, where

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

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

$$W_p = \{x | d(p,x) \geqslant d(q,x), p \neq q \in S\}, \qquad p \in S.$$

We will need the following properties of these diagrams.

Fact 1:

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

Fact 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:

The regions V_p and W_p are unbounded if and only if p is a point on the convex hull of S. Otherwise V_p is a nonempty convex polygon and W_p is empty.

Two points p and q of S are said to be *closest* (respectively *furthest*) *point Voronoi neighbors* if V_p and V_q (resp. W_p and W_q) share a common point.

Fact 4:

Two points p and q of S are closest *and* furthest point Voronoi neighbors if and only if (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 if and only if they are closest (resp. furthest) point Voronoi neighbors.

Fact 5:

Both the closest and furthest point Voronoi diagram (as well as the respective Delaunay triangulations) of n points can be constructed in $O(n \log n)$ 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 1 to 5 and other properties of these constructions consult [17]. 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 1: The α-shape of S is a subgraph of $DT_f(S)$ if $\alpha \geqslant 0$ and a subgraph of $DT_c(S)$ if $\alpha \leqslant 0$.

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-neighbors, then they are adjacent in the respective Delaunay triangulation. We consider three cases:

a) $\alpha = 0$: The convex hull is a subgraph of both $DT_c(S)$ and $DT_f(S)$ by Fact 4.
b) $\alpha > 0$: Let p and q be α-neighbors and let c be the center of the disc of radius $1/\alpha$ that touches p and q and contains all other points $r \in S$. Clearly $d(c,p) \geqslant d(c,r)$ and $d(c,q) \geqslant d(c,r)$ for all $r \in S$. Thus c is contained in both W_p and W_q and hence p and q must be furthest point Voronoi neighbors. Therefore p and q are adjacent in $DT_f(S)$.
c) $\alpha < 0$: The proof is essentially the same as in b) replacing W by V, furthest by closest, contains all by contains no, $\geqslant$ by $\leqslant$, 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.

Lemma 2: For each point $p \in S$ there exists a real number $\alpha_{\max}(p)$ such that p is α-extreme in S if and only if $\alpha \leqslant \alpha_{\max}(p)$.

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

a) *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 $r \geqslant d(p, W_p) = \min\{d(p,x) | x \in W_p\}$. Thus p is α-extreme for $0 < \alpha \leqslant 1/d(p, W_p)$. A convex hull point is trivially α-extreme for nonpositive α, hence p is α-extreme for all $\alpha \leqslant 1/d(p, W_p) = \alpha_{\max}(p)$.
b) *p is not a convex hull point of S*: it is easy to see that p cannot be α-extreme for $\alpha \geqslant 0$. 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 2 and 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 $r \leqslant d_p = \max\{d(p, x) | x \in V_p\}$. This implies that p is α-extreme for all $\alpha \leqslant -1/d_p = \alpha_{\max}(p)$. Q.E.D.

Note that for a point $p \in S$ the vertices of the polygons V_p (resp. W_p) are the centers of the circumscribed circles of the triangles in $DT_c(S)$ (resp. $DT_f(S)$) which involve p. Clearly for $p \in S$ the vertices of V_p or the vertices and edges of W_p are sufficient to determine $\alpha_{\max}(p)$. Thus it is not difficult to see how the set $\{\alpha_{\max}(p) | p \in S\}$ can be computed in linear time given $DT_c(S)$ and $DT_f(S)$.

Lemma 3: For every edge e belonging to either $DT_c(S)$ or $DT_f(S)$ there are real numbers $\alpha_{\min}(e)$ and $\alpha_{\max}(e)$, with $\alpha_{\min}(e) \leqslant \alpha_{\max}(e)$, such that e is an edge of the α-shape of S if and only if $\alpha_{\min}(e) \leqslant \alpha \leqslant \alpha_{\max}(e)$.

Proof: First we state without proof the following two facts.

Fact 6:

Given a point p and a semi-infinite line segment s there exists a positive real number $a = a(p, s)$ such that $\{d(p, x) | x \in s\} = [a, \infty)$.

Fact 7:

Given a point p and a closed line segment s there exist positive real numbers $a = a(p, s)$ and $b = b(p, s)$ such that $\{d(p, x) | x \in s\} = [a, b]$.

Now, let p and q be the two endpoints of an edge e. We have to consider three cases.

a) *e is an edge of $DT_c(S)$ but is not a convex hull edge*: 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, which we call v, bounding both of the Voronoi regions V_p and V_q (i.e., v is the Voronoi edge dual to e). Thus by Fact 7 there are discs of radius r touching p and q and not containing other points of S for exactly those r satisfying, $a(p, v) \leqslant r \leqslant b(p, v)$. It is easy to see that as a consequence of Fact 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 for exactly those α satisfying, $-1/a(p, v) = \alpha_{\min}(e) \leqslant \alpha \leqslant \alpha_{\max}(e) = -1/b(p, v)$.

b) *e is an edge of $DT_f(S)$ but is not a convex hull edge*: the proof is essentially the same as in a) replacing "furthest" by "closest," "contains all" by "contains no," etc.

c) *e is a convex hull edge*: first note that p and q are trivially α-neighbors for $\alpha = 0$. 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 bounding both of the furthest point Voronoi regions W_p and W_q. Thus by Fact 6 there are discs of radius r touching p and q and containing all other points of S for exactly those r satisfying $r \geqslant a(p, w)$.

By the same argument there are discs of radius r touching p and q and containing none of the other points of S for exactly those r satisfying $r \geqslant a(p, v)$, where v is the semi-infinite line segment bounding V_p and V_q. Thus p and q are α-neighbors, i.e., e is and edge of the α-shape, for all α satisfying $-1/a(p, v) = \alpha_{\min}(e) \leqslant \alpha \leqslant \alpha_{\max}(e) = 1/a(p, w)$. Q.E.D.

Note that the line segment which borders both V_p and V_q (resp. W_p and W_q) of two Voronoi neighbors p and q is contained in the perpendicular bisector between p and q and that its endpoint(s) is (are) the center(s) of the circumscribed circle(s) of the (at most two) triangle(s) of $DT_c(S)$ (resp. $DT_f(S)$) which involve both p and q. Thus given $DT_c(S)$ and $DT_f(S)$, $\alpha_{\min}(e)$ and $\alpha_{\max}(e)$ can clearly be computed in constant time for each edge e.

IV. Construction of α-Shapes and the Shape Spectrum

A. α-Shapes

Together, Lemmas 1, 2, and 3 give rise to the following algorithm for determining the α-shape of a set S.

Algorithm 1 (Construction of the α-shape of S).

1) Construct DT:
   if $\alpha \geqslant 0$, 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 1 and 2.
3) Determine the α-neighbors of S.
   Again, DT contains all the information necessary to perform this task; see also Lemma 1 and 3.
4) Output the α-shape.
   Output the graph on the α-extreme points with all α-neighbor connections.

The correctness of Algorithm 1 follows immediately from Lemmas 1, 2, and 3. A straightforward analysis of Algorithm 1 gives rise to the following.

Theorem 1: The α-shape of n points in the plane can be determined for an arbitrary real α in time $O(n \log n)$ and space $O(n)$.

Proof: It suffices to show that the stated bounds hold for Algorithm 1. Step 1) can be done in $O(n \log n)$ time and $O(n)$ space by Fact 5. Once the appropriate Delaunay triangulation has been constructed Steps 2), 3), and 4) can be done (see the notes following Lemmas 2 and 3) in $O(n)$ time and $O(n)$ space. Whenever in Step 2) or 3) the actual value of $\alpha_{\max}(p)$, or $\alpha_{\min}(e)$, or $\alpha_{\max}(e)$ 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.

B. Interior and Exterior Faces

It should be clear that Algorithm 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 bounding 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 noncontinuous 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 $\alpha \geqslant 0$ (resp. $\alpha \leqslant 0$) 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.

C. 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 1 makes clear, it is possible to construct α-shapes, following an initial expenditure of $O(n \log n)$ to construct both Delaunay triangulations, 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 7: The *shape spectrum* $SP(S)$ of a point set S is defined to be the set of intervals $\text{int}(p) = (-\infty, \alpha_{\max}(p)]$ and $\text{int}(e) = [\alpha_{\min}(e), \alpha_{\max}(e)]$, $p \in S$, 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: The shape spectrum $SP(S)$ of a set S of n points can be constructed in time $O(n \log n)$ and space $O(n)$.

Proof: Immediate generalization of Algorithm 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_0, find the α_0-shape of S. This can be done trivially in linear time by determining all points p and edges e such that $\alpha_0 \in \text{int}(p)$ and $\alpha_0 \in \text{int}(e)$. However, by using a more advanced data structure to store the intervals of $SP(S)$, such as Edelsbrunner's rectangle tree [5], called a tile tree in the independent paper of McCreight [13], the α_0-shape of S can actually be constructed in time $O(\log n + t)$, where t is the number of points and edges in the α_0-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_0 such that the α_0-shape of S contains exactly k points. If the endpoints of the intervals $\text{int}(p)$, $p \in S$, are stored in a sorted array, α_0 can clearly be found in constant time.

3) A similar problem addresses the fine tuning of α-shapes: given the α_0-shape of S for some α_0, find the largest $\alpha_1 < \alpha_0$ (or the smallest $\alpha_1 > \alpha_0$), such that the α_1-shape is different from the α_0-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, $\text{int}(e)$ 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 $\text{int}(x) \in SP(S)$, 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 $\text{int}(x) \in SP(S)$, 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.

V. Constructing the α-Hull

In the preceding section we presented an $O(n \log n)$ 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 $O(n \log n)$ time. As these facts seem to be quite obvious for $\alpha \geqslant 0$, 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 $r \leqslant R$, 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 $B(x, r)$ denote the open disc of radius r centered at x.

Lemma 5: 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 $D = B(c, r)$ be a disc which does not contain any points of S. Let $p \in S$ be the point, such that $d(c, p) = \min\{d(c, x) | x \in S\}$. Clearly the disc $D' = B(c, d(c, p))$ 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 $D_1 = B(t, d(t, p))$ 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 $D \subset D' \subset D_1$ (Fig. 5). Q.E.D.

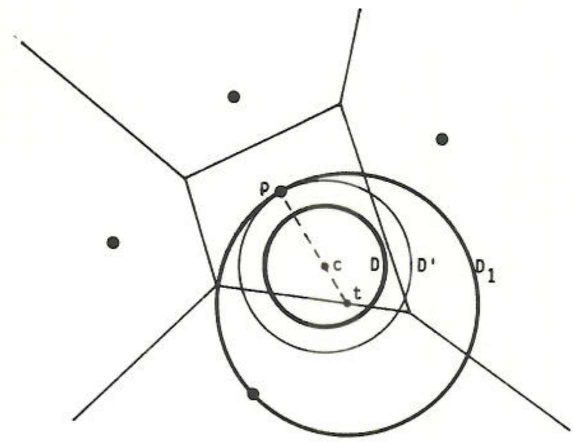

Fig. 5. Nested empty discs D, D', D_1.

As a consequence of Lemma 5 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 on 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 8:

> 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 (Fig. 6). Then
> $$\cup \{B(x, d(x, p)) | x \in L\} = B(a, d(a, p)) \cup B(b, d(b, p)).$$

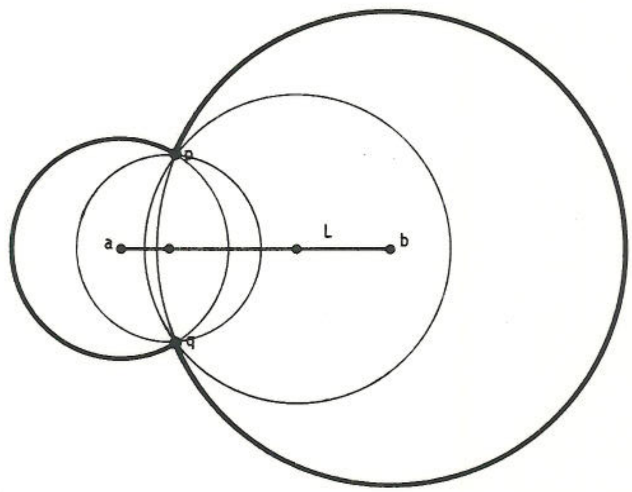

Fig. 6. Centers on a bounded line segment.

Fact 9:

> 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
> $$\cup \{B(x, d(x, p)) | x \in L\} = B(a, d(a, p)) \cup H(p, q),$$
> where $H(p, q)$ denotes the open halfplane defined by the straight line through p and q which contains the infinite portion of L (Fig. 7).

Lemma 6: 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 we only have to consider appropriate discs centered on edges of $VD_c(S)$. Let $p, q \in S$ be two Voronoi neighbors and let r be the edge bounding both V_p and V_q. Clearly for every $x \in r$,

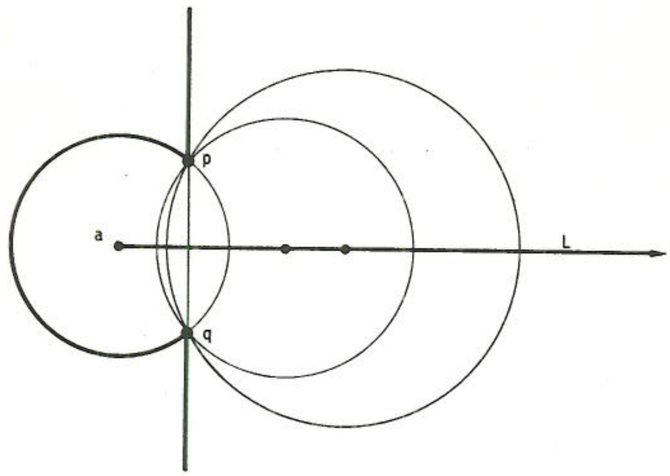

Fig. 7. Centers on a semi-infinite line segment.

$B(x, d(x, p))$ contains no points of S. Now observe that the set $\{x \in r | d(x, p) \geqslant -1/\alpha\}$ is either empty or forms one or two closed line segments. (If p and q are convex hull points, then one line segment is semi-infinite.) Thus by Facts 8 and 9 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 four 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 theorem.

Theorem 2: The α-hull of a set S of n points can be computed in time $O(n \log n)$ using $O(n)$ space.

Proof: We consider two cases:

a) $\alpha < 0$: By Lemma 6 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 $O(n \log n)$ 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 1.

b) $\alpha \geqslant 0$: For nonnegative 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 Jarvis' [9], may be preferred. The details are left to the reader. Q.E.D.

Conclusion

In this paper we have developed the notion of the α-hull and α-shape of a set of points in the plane. We presented efficient algorithms for constructing 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 believe that α-shapes will be very useful 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 rectangle tree or the sorted lists) which were mentioned in Section IV.

Finally, we want to point out that the notion of α-shapes generalizes nicely to point sets in 3 or k dimensions. One can define α-extreme points, α-neighbors and α-triples, and so on, in a manner similar to the definitions of Section II, by using balls of radius $1/\alpha$ instead of discs. The 3-dimensional α-shape is related to 3-dimensional Voronoi diagrams in a way similar to the relationship between planar α-shapes and planar Voronoi diagrams. Lemmas 1 to 3 and Algorithm 1 carry over to 3 dimensions without much modification. Applying the results of Seidel [16] for finding the 3-dimensional Voronoi diagram, the 3-dimensional α-shape can be constructed on $O(n^2)$ time. As the Voronoi diagram and therefore the α-shape in 3 dimensions has quadratic space-complexity in the worst case this is also optimal.

Acknowledgment

We wish to express our sincere appreciation for the numerous helpful comments made by the referees.

References

[1] S. G. Akl, and G. T. Toussaint, "Efficient convex hull algorithms for pattern recognition applications," in *Proc. 4th Intern. Joint. Conf. on Pattern Recognition*, Kyoto, pp. 483–487, 1978.

[2] K. Q. Brown, "Voronoi diagrams from convex hulls," *Info. Proc. Letters*, vol. 9, pp. 223–228, 1979.

[3] ——, "Geometric transforms for fast geometric algorithms," Rep. CMU-CS-80-101, Carnegie-Mellon Univ., Dept. of Comp. Sci., 1980.

[4] A. K. Dewdney and J. K. Vranch, "A convex partition of R^3 with applications to Crum's problem and Knuth's post office problem," *Utilitas Math.*, vol. 12, pp. 193–199, 1977.

[5] H. Edelsbrunner, "A new approach to intersections: Part I," to appear in *Intern. J. Comp. Math.*, 1983.

[6] J. Fairfield, "Contoured shape generation: Forms that people see in dot patterns," in *Proc. IEEE Conf. on Cybernetics and Society*, pp. 60–64, 1979.

[7] I. G. Gowda and D. G. Kirkpatrick, "Exploiting linear merging and

extra storage in the maintenance of fully dynamic geometric data structures," in *Proc. 18th Annual Allerton Conf. on Commun., Contr., and Comput.*, pp. 1–10, Oct. 1980.
[8] R. L. Graham, "An efficient algorithm for determining the convex hull of a finite planar set," *Info. Proc. Letters*, vol. 1, pp. 132–133, 1972.
[9] R. A. Jarvis, "Computing the shape hull of points in the plane," in *Proc. of the IEEE Comp. Soc. Conf. on Pattern Recognition and Image Processing*, pp. 231–241, 1977.
[10] H. W. E. Jung, "Ueber die kleinste Kugel, die eine raeumlich Figur einschliesst," *J. f.d. Reine u. Angew. Math.*, vol. 123, pp. 241–257, 1901.
[11] V. E. Klee, "On the complexity of *d*-dimensional Voronoi diagrams," *Archiv d. Math.*, vol. 34, pp. 75–80, 1980.
[12] D. W. Matula and R. R. Sokal, "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] E. M. McCreight, "Efficient algorithms for enumerating intersecting intervals and rectangles," *J. Algorithms*, vol. 2, pp. 245–260, 1981.
[14] M. H. Overmars, "Dynamization of order decomposable set problems," Rep. RUU-CS-80-9, Univ. of Utrecht, Dept. of Comp. Sci., 1980.
[15] F. P. Preparata, "Steps into computational geometry," Rep. R-792, Univ. of Illinois, Coordinated Sci. Lab., 1977.
[16] R. Seidel, "A convex hull algorithm optimal for point sets in even dimensions," Report 81-14, University of British Columbia, Dept. of Comp. Sci., 1981.
[17] M. I. Shamos and D. Hoey, "Closest point problems," in *Proc. 16th IEEE Symp. on Foundations of Comput. Sci.*, pp. 151–162, 1975.
[18] K. J. Supowit, "The relative neighborhood graph with an application to minimum spanning trees," Tech. Rep., Univ. of Illinois, Dept. of Comput. Sci., Aug. 1980.
[19] G. T. Toussiant, "The relative neighborhood graph of a finite planar set," *Patt. Recog.*, vol. 12, pp. 261–268, 1980.
[20] ——, "Pattern recognition and geometrical complexity," *5th International Conference on Pattern Recognition*, Miami Beach, FL, pp. 1324–1347, Dec. 1980.

SECT—A Coding Technique for Black/White Graphics

HIROYOSHI MORITA AND SUGURU ARIMOTO, FELLOW, IEEE

Abstract—**A new coding technique called SECT for digitized two-tone or black / white pictures which is particularly applicable to polygonal objects composed of horizontal, vertical, and diagonal lines such as logical circuit patterns, characteristic font patterns, and mechanical drawings is described. While many previously proposed coding techniques attempt to encode the positions of transitive elements in a picture by using codes such as Huffman codes or Wyle codes which are constructed on the basis of picture statistics, selective element coding technique (SECT) focuses its attention on reducing the number of transitive elements in each picture independently of the statistics of the picture ensemble. The relation of the number of selective elements and objects in the picture is discussed. Furthermore decoding algorithm to reproduce a picture from selective elements and its decodability are described.**

I. Introduction

A NUMBER of coding techniques for digitized black/white pictures have been proposed, and their data compression performance has been studied by computer simulations for pictures such as graphics, documents, weather maps, and Chinese character patterns [1], [2], [3]. In this paper we consider a two-dimensional coding technique for black/white graphics which is particularly applicable for polygonal objects composed of horizontal, vertical, and diagonal lines, e.g., logical circuit patterns generated by CAD systems, characteristic font patterns, and mechanical drawings.

Manuscript received August 10, 1981; revised December 10, 1982.
The authors are with Faculty of Engineering Science, Osaka University, Toyonaka, Osaka, Japan.

Before describing our proposed coding technique, it is worthwhile to review previous coding techniques. Considerable work has been done in digital black/white picture coding for facsimile transmission. Run-length coding [4], prediction coding [5], [6], [7], and READ coding [8] are well-known, and these have been realized as codes in actual facsimile equipment. In these coding schemes, the encoding process performs by scanning a picture sequentially line by line from top to bottom. The strategy of run-length coding is to encode the length of each black or white run by means of Huffman codes [9] or Wyle codes [10]. The idea of prediction coding such as the well-known TUH code [7] is to predict each picture element by using some function of neighboring elements and to encode the resulting error sequence by run-length coding. READ coding, as well as RAC [11] and EDIC [12], encodes the distance between the start position of a run and that of a certain previously scanned run which satisfies some prescribed conditions.

These coding schemes contain two common operations. The first operation is to select out the start position of each run from a given picture, where the start position of a run is often called a horizontal transitive element or simply a transitive element. A transitive element is formally defined as one whose value is different from the previous one in a horizontal direction. The second operation is to memorize the positions of transitive elements in some way, e.g., the distance between two consecutive transitive elements, or equivalently, the length of a run. Data compression is

0018-9448/83/0700-0559$01.00