ELSEVIER

Discrete Applied Mathematics 122 (2002) 37–54

DISCRETE APPLIED MATHEMATICS

2-Point site Voronoi diagrams ☆

Gill Barequet[a,*,1] Matthew T. Dickerson[b,2], Robert L. Scot Drysdale[c,2]

[a] *Faculty of Computer Science, The Technion—Israel Institute of Technology, Haifa 32000, Israel*
[b] *Department of Mathematics and Computer Science, Middlebury College, Middlebury, VT 05753, USA*
[c] *Department of Computer Science, Dartmouth College, 6211 Sudikoff Laboratory, Hanover, NH 03755-3510, USA*

Received 9 July 1999; received in revised form 16 March 2001; accepted 30 July 2001

Abstract

In this paper, we define a new type of a planar distance function from a point to a *pair* of points. We focus on a few such distance functions, analyze the structure and complexity of the corresponding nearest- and furthest-neighbor Voronoi diagrams (in which every region is defined by a pair of point sites), and show how to compute the diagrams efficiently. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Distance function; Voronoi diagram; Lower envelope; Davenport–Schinzel theory

1. Introduction

The standard Voronoi diagram of a set of n given points (called sites) is a subdivision of the plane into n regions, one associated with each site. Each site's region consists of all points in the plane closer to it than to any of the other sites. One application that frequently occurs is what Knuth called the "post office" problem. Given a letter to be delivered, the nearest post office to the destination can be found by locating

☆ A preliminary version of this paper appeared in WADS '99

* Corresponding author.
E-mail addresses: barequet@cs.technion.ac.il (G. Barequet), dickerso@middlebury.edu (M.T. Dickerson), scot@cs.dartmouth.ed (R.L. Scot Drysdale).

[1] Supported in part by the U.S. Army Research Office under Grant DAAH04-96-1-0013. The work was done while this author was affiliated with the Department of Computer Science at Johns Hopkins University, Baltimore, MD, USA.
[2] Supported in part by the National Science Foundation under Grant CCR-93-1714.

PII: S0166-218X(01)00320-1

the destination point in the Voronoi diagram of the post office sites. This is called a "locus approach" to solving the problem—points in the plane are broken into sets by the answer to a query (in this case, "Which post office is nearest?"). All points that give the same answer are in the same set. Answering queries is reduced to planar point location once the nearest-neighbor diagram is computed. When the furthest site is sought for every point in the plane, we obtain the furthest-neighbor Voronoi diagram.

The Voronoi diagram has been rediscovered many times in dozens of fields of study including crystallography, geography, metrology, and biology, as well as mathematics and computer science. A comprehensive review of the various variations of Voronoi diagrams and of the hundreds of applications of them is given by Okabe et al. [9].

In particular, there have been a number of studies of variants of the Voronoi diagrams based on nonEuclidean distance functions and on sites that are line segments, circles, polygons, and other shapes more complicated than points. Another studied variant is the kth-order Voronoi diagram. Here the plane is broken into regions where all points in a given region have the same k sites as their k nearest neighbors. However, even in this case the distance measure is based only on the pairwise distance.

The regular (1-site) nearest-neighbor Voronoi diagram (with respect to the Euclidean distance function) can be viewed as the result of blowing circles around each point, where each point in the plane belongs to the region of the site whose circle sweeps it first. (Similarly, the furthest-neighbor diagram is constructed by considering, for each point in the plane, the last circle that sweeps it.) Note that: 1. All the circles start to grow at the *same* "time" $t = 0$ (representing the zero distance from the sites); and 2. All the circles grow in the same speed. 2-site Voronoi diagram are the results of blowing some family of shapes around each *pair* of sites. Each 2-site distance function is modeled by a different blown shape, and has a different setting of the initial times and growing rates of the respective shapes.

It is well known that the 1-site nearest- (resp., furthest-) neighbor Voronoi diagram is the xy-projection of the lower (resp., upper) envelope of the xy-monotone surfaces modeling the functions that measure the distance from each site. For every location (x, y), the z coordinate of the respective Voronoi surface of the site p is the two-dimensional distance from (x, y) to p. For the Euclidean distance function all the surfaces are copies of the same cone whose apex is translated to the sites. For 2-site distance functions, each pair of sites (p, q) is associated with a surface, where for every point (x, y, z) on the surface, the value z is the 2-site distance from (x, y) to the pair (p, q). For each 2-site distance function we describe the associated family of Voronoi surfaces, the xy-projection of whose envelopes form the respective Voronoi diagrams.

The paper is organized as follows. In the rest of this section, we motivate this study, introduce several 2-site distance functions, and summarize our results on the nearest- and furthest-neighbor Voronoi diagrams that are based on these distance functions. In Section 2, we study the sum- and product-of-distances distance functions, whose respective diagrams turn out to be well known but under a different definition. In Sections 3–5, we investigate the respective Voronoi diagrams of the triangle area, distance from a line (or a segment), and difference-between-distances distance functions. We discuss the properties of each distance function, give bounds on the complexities

of the respective diagrams, and describe efficient algorithms for computing them. We terminate in Section 6 with some concluding remarks.

1.1. Motivation and applications

Our study of 2-point site distance functions was motivated by the famous Heilbronn's triangle problem:

> Let $\{P_1, P_2, \ldots, P_n\}$ be a set of n points in $[0,1]^2$, such that the minimum of the areas of the triangles $P_iP_jP_k$ (for $1 \leqslant i < j < k \leqslant n$) assumes its maximum possible value $\mathscr{H}(n)$. Estimate $\mathscr{H}(n)$.

Assume that Δ is an estimate of the solution of this problem. Then, each pair of points defines a strip, centered at the line joining the two points, that cannot contain any other point of the set. The width ω of this strip is inversely-proportional to the distance d between the two points. (Specifically, $\omega = 4\Delta/d$.) Imagine now varying Δ continuously from 0 up. This will "pump up" the forbidden strips in the same way the strips of the $\mathscr{A}$ (triangle area) distance function are growing. Showing that the strips cannot grow too much, for otherwise at least one of them would contain another point of the set, is a known technique for setting an upper bound on Δ.

For finding applications of the other 2-site distance functions, we may consider them as representing some "cost" of placing an object at a point v with respect to two reference points p and q. For example, The function $\mathscr{S}$ (sum of distances) can be regarded as a variant of the post-office problem, in which one needs to send a letter from *two* different post offices, so that the receiver will be able to compare the two arriving copies of the message and verify its correctness. The function $\mathscr{A}$ (triangle area) can model two envoys sent from v to p and q, where this time the envoys maintain a live connection between them, so that the cost is the area swept in between the two paths. The function $\mathscr{D}$ (difference between distances) can measure the quality of a stereo sound, where speakers are positioned at the sites.

1.2. Our results

In contrast with the 1-site distance functions studied so far, we define several distance functions from a point to a *pair* of points in the plane. We denote by $d(a,b)$ the Euclidean distance between the points a and b, and by $A(a,b,c)$ the area of the triangle defined by the points a, b, and c. For two points a and b, we denote by ℓ_{ab} the line defined by a and b, and by $\overline{ab}$ the line segment whose endpoints are a and b. Given two point sites p and q, we define the following distance functions from a point v to the pair (p,q):

(1) Sum of distances: $\mathscr{S}(v,(p,q)) = d(v,p) + d(v,q)$; and product of distances: $\mathscr{M}(v,(p,q)) = d(v,p) \cdot d(v,q)$.
(2) Triangle area: $\mathscr{A}(v,(p,q)) = A(v,p,q)$.

Table 1
Worst-case combinatorial complexities of $V_{\mathscr{F}}^{(n|f)}(S)$

$\mathscr{F}$	$\mathscr{S}$, $\mathscr{M}$	$\mathscr{A}$	$\mathscr{L}$	$\mathscr{G}$	$\mathscr{D}$
Nearest-neighbor diagram	$\Theta(n)$	$\Theta(n^4)$	$\Theta(n^4)$	$\Theta(n^4)$	$\Omega(n^4)$, $O(n^{4+\varepsilon})$
Furthest-neighbor diagram	$\Theta(n)$	$\Theta(n^2)$	$\Theta(n^2)$	$\Theta(n)$	$\Theta(n^2)$

(3) Distance from a line: $\mathscr{L}(v,(p,q))=\min_{u\in\ell_{pq}} d(v,u)$; and Distance from a segment: $\mathscr{G}(v,(p,q))=\min_{u\in\overline{pq}} d(v,u)$.
(4) Difference between distances: $\mathscr{D}(v,(p,q))=|d(v,p)-d(v,q)|$.

All these 2-site distance functions are symmetric in p and q. (Some of them are symmetric in all of v, p, and q, but this has no importance here.) All these functions, like the regular Euclidean distance function, are invariant under translations and rotations of the plane. For every 2-site distance function $\mathscr{F}$ we define the nearest- (resp., furthest-) neighbor Voronoi diagram (with respect to $\mathscr{F}$) of a point set S as the partition of the plane into regions, each corresponding to a pair of points of S. We denote these diagrams by $V_{\mathscr{F}}^{(n|f)}(S)$. The region that corresponds to $p,q\in S$ consists of all the points $v\in\Re^2$ for which $\mathscr{F}(v,(p,q))$ is minimized (or maximized), where the optimum is taken over all the pairs of points in S. We denote by "cells" the connected components of the diagram. (A region may consist of multiple cells.) We summarize in Table 1 the major results of this paper: the bounds on the largest diagram complexities for these distance functions. (For some distance functions there exist point sets whose respective Voronoi diagram has complexity less than that of the worst case.)

In the next sections, we analyze the nearest- and furthest-site Voronoi diagrams of point sets in the plane with respect to the distance functions defined above.

2. Sum and product of distances

We begin with the simple sum- and product-of-distances distance functions:

Definition 1. Given two points p,q in the plane, the "distances" $\mathscr{S}(v,(p,q))$ and $\mathscr{M}(v,(p,q))$ from a point v in the plane to the unordered pair (p,q) are defined as $d(v,p)+d(v,q)$ and $d(v,p)\cdot d(v,q)$, respectively.

Given a set S of n points in the plane, we wish to find its nearest-neighbor Voronoi diagram with respect to the distance function $\mathscr{S}$ (or $\mathscr{M}$). In this case, it is the partition of the plane into regions for each pair of sites $p,q\in S$, which contain all points in the plane closer to (p,q) than to any other pair of sites in S.

We start by describing the surfaces (or blown shapes) which describe these distance functions. The curve $\mathscr{S}(v,(p,q))=c$ (for a fixed pair of points p and q and a constant $c\geqslant d(p,q)$) is an ellipse. (For $c<d(p,q)$ the curve is empty, and for $c=d(p,q)$ the ellipse degenerates to a segment.) Thus, $V_{\mathscr{S}}^{(n)}(S)$ can be viewed as the result of blowing ellipses around each pair of points of S, so that the two points remain the

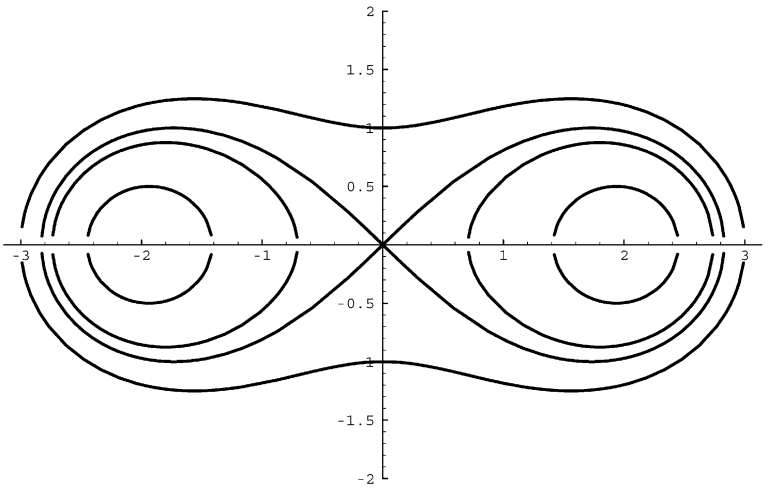

Fig. 1. $\mathscr{M}(v,(p,q))=c$ for $p=(-2,0)$, $q=(2,0)$, and $c=2,3.5,4,5$.

foci of the blown ellipse. Every point v in the plane belongs to the region of the pair of sites whose ellipse sweeps it first. In this construction, the ellipses starts to grow at *different* times: the initial ellipse (the analog of the 0-radius circle) whose foci are the sites p and q is the straight line-segment that connects between p and q. The ellipse start to grow at time $t=d(p,q)$, since for every point v on that segment $\mathscr{S}(v,(p,q))=d(p,q)$.

The curve $\mathscr{M}(v,(p,q))=c$ (for fixed points p,q and $c>0$) is the oval of Cassini [12, p. 44]. For $0<c\leqslant d^2(p,q)/4$ the curve consists of two "leaves" drawn around p and q and symmetric around the bisector of the line-segment $\overline{pq}$. (See Fig. 1 for an illustration.) When $c=d^2(p,q)/4$ the leaves touch at the midpoint of $\overline{pq}$ and form a lemniscate [12, p. 40]. For $c>d^2(p,q)/4$ the two leaves are merged into one quartic curve.

Though the blown shapes for both distance functions are very different from the growing circles of the Euclidean metric, it turns out that the respective Voronoi diagrams are closely related:

Fact 2. *Let $v\notin S$ be a point in the plane. Also, let $(p,q)_{\mathscr{S}}$ and $(p,q)_{\mathscr{M}}$ be the closest pairs of points of S to v according to $\mathscr{S}$ and $\mathscr{M}$, respectively, and let p' and q' be the two closest sites of S to v with respect to the regular Euclidean distance function. Then the (unordered) pairs $(p,q)_{\mathscr{S}}$, $(p,q)_{\mathscr{M}}$ and (p',q') are all identical.*

This simply tells us that $V^{(n)}_{\mathscr{S}}(S)$ and $V^{(n)}_{\mathscr{M}}(S)$ are identical to the second-order nearest-neighbor Voronoi diagram of S with respect to the regular Euclidean distance function. (With the only difference that the points of S are singular points in $V^{(n)}_{\mathscr{M}}(S)$, since $\mathscr{M}(p,(p,q))=0$ for all $q\neq p$ in S, hence the points of S are isolated vertices of $V^{(n)}_{\mathscr{M}}(S)$ in which no Voronoi edge occurs.) It is well known that the edges of this

diagram (portions of bisectors of pairs of point sites) are straight line-segments. This may seem at first surprising, since the bisectors between the regions of two pairs of sites (p,q) and (r,s) (for both $\mathscr{S}$ and $\mathscr{M}$) are in general much more complex curves. The reason for this is that the Voronoi diagram contains only portions of bisectors of pairs which share one site, that is, of the form (p,q) and (p,r). The combinatorial complexity of the second-order (Euclidean) Voronoi diagram is known to be $\Theta(n)$ [10,6]. The diagram can be computed in optimal $\Theta(n\log n)$ time and $\Theta(n)$ space.

Similarly, the diagrams $V^{(f)}_{\mathscr{S}}(S)$ and $V^{(f)}_{\mathscr{M}}(S)$ are identical to the second-order furthest-neighbor Voronoi diagram of S with respect to the regular Euclidean distance function. The bounds on the complexity of the diagram and on the time needed to compute it are the same as for the nearest-neighbor diagram.

3. Triangle area

3.1. Growing strips

We now define the 2-site triangle-area distance function:

Definition 3. Given two points p,q in the plane, the "area-distance" $\mathscr{A}(v,(p,q))$ from a point v in the plane to the unordered pair (p,q) is defined as $A(v,p,q)$, the area of the triangle defined by the three points.

For a fixed pair of points p and q, the curve $\mathscr{A}(v,(p,q))=c$, for a constant $c\geqslant 0$, is a pair of parallel lines at distance $4c/d(p,q)$ apart. The respective 2-site Voronoi diagram is constructed by blowing infinite strips, each strip centered at the line ℓ_{pq} passing through the points $p,q\in S$. All the strips start to grow simultaneously at $t=0$, however at different rates. The growing rate of the strip defined by the points p and q is inversely proportional to $d(p,q)$. Each point v in the plane belongs to the region of the pair of sites whose strip sweeps it first (or last).

The bisector between the regions of two pairs of points (p,q) and (r,s) is a pair of straight lines passing through the intersection point of ℓ_{pq} and ℓ_{rs}. Figs. 2(a) and (b)

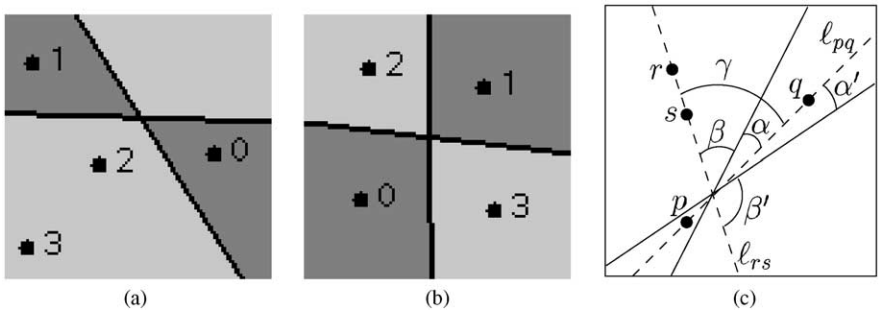

(a) (b) (c)

Fig. 2. Bisectors for $\mathscr{A}$.

show two configurations of four points. The region of points closer to $(0,1)$ (according to $\mathscr{A}$) is shown in dark gray, while the region of $(2,3)$ is shown in light gray. The slopes of the bisector lines are weighted averages of the slopes of ℓ_{pq} and ℓ_{rs} (see Fig. 2(c)). Let α (resp., β) be the angle between a bisector line and ℓ_{pq} (resp., ℓ_{rs}), and let $\gamma = \alpha + \beta$ be the angle between ℓ_{pq} and ℓ_{rs}. It is easily seen that $\sin\alpha/\sin\beta = d(r,s)/d(p,q)$. A simple calculation shows that

$$\tan\alpha = \frac{d(r,s)\sin\gamma}{d(p,q) + d(r,s)\cos\gamma}.$$

Similarly,

$$\tan\alpha' = \frac{d(r,s)\sin(\pi-\gamma)}{d(p,q) + d(r,s)\cos(\pi-\gamma)} = \frac{d(r,s)\sin\gamma}{d(p,q) - d(r,s)\cos\gamma}.$$

3.2. Nearest-neighbor diagram

We first lower bound the complexity $V^{(n)}_{\mathscr{A}}(S)$. The $\Theta(n^2)$ strips defined by all the pairs of points start growing at the same "time," thus their respective regions are not empty, and the intersection point of each pair of such zero-width strips is a feature of the diagram. It follows that the combinatorial complexity of $V^{(n)}_{\mathscr{A}}(S)$ is $\Omega(n^4)$.

We now upper bound the complexity of the diagram. Refer to the growing strip that corresponds to two point sites $p = (p_x, p_y)$ and $q = (q_x, q_y)$. The Voronoi surface that corresponds to p and q consists of a pair of halfplanes, both bounded by the line ℓ_{pq}, and ascending outward of it. The slope of the two halfplanes is $1/2d(p,q)$. More precisely, the surface that corresponds to p and q is the bivariate function of a point $v = (v_x, v_y)$

$$F^{p,q}_{\mathscr{A}}(v_x, v_y) = A(v,p,q) = \frac{1}{2}\mathrm{abs}\left(\begin{vmatrix} v_x & v_y & 1 \\ p_x & p_y & 1 \\ q_x & q_y & 1 \end{vmatrix}\right).$$

Since the surfaces are piecewise-linear bivariate functions in $\Re^3$, we can apply Theorem 7.1 of [11, p. 179] and obtain a slightly super-quartic upper bound, namely, $O(n^4\alpha(n^2))$,[3] on the complexity of $V^{(n)}_{\mathscr{A}}(S)$. However, we can do better than that. The complexity of the diagram is no more than the complexity of the *zone* of the plane $z = 0$ in the arrangement of the planes obtained by extending the $\Theta(n^2)$ halfplanes mentioned above. The latter complexity is $\Theta(n^4)$ in the worst case [3] (see also [11, p. 231, Theorem 7.50]).

Hence we have the following:

Theorem 4. *The combinatorial complexity of* $V^{(n)}_{\mathscr{A}}(S)$ *is* $\Theta(n^4)$.

[3] $\alpha(n)$ is an extremely slowly growing functional inverse of Ackermann's function. For all practical values of n, $\alpha(n)$ does not exceed a very small constant.

Note that this bound applies for *all* diagrams of sites in general position (no three collinear sites, and no point common to three lines defined by the sites) and not only for the worst case.

Fig. 3(a) shows three points in the plane, while Fig. 3(b) shows the three respective Voronoi surfaces in a perspective view. Fig. 3(c) shows the same construction from below. Fig. 3(d) shows the Voronoi diagram of the three points, which is the xy-projection of the lower envelope of the surfaces.

In fact, a single pair of sites can have $\Theta(n^4)$ cells in the diagram. To see this, put two points p and q *very* close together and spread all the other points far apart. The $\Theta(n^2)$ lines passing through all pairs of points form an arrangement with $\Theta(n^4)$ faces, in which cells begin to "grow" as the strips expand. We can position p and q close enough to make their respective strip grow fast enough so as to "bypass" all the other strips and grab a piece of each face of the arrangement.

3.3. Furthest-neighbor diagram

Theorem 5. *The combinatorial complexity of* $V^{(f)}_{\mathscr{A}}(S)$ is $\Theta(n^2)$ *in the worst case.*

Proof. The lower bound is set by an example. Let n be divisible by 4. Put n points evenly spaced around the unit circle C.

Lemma 6. *The center* o *of* C *is a vertex in* $V^{(f)}_{\mathscr{A}}(S)$, *shared by the regions of all the pairs of points at distance* $\pi/2$ *along* C.

Proof. Let $p_0=(1,0)$ be the first point, and order the points counterclockwise along C, so that $p_i=(\cos(2i\pi/n),\sin(2i\pi/n))$. A simple calculation shows that $\mathscr{A}(o,(p_0,p_i))=\sin(2i\pi/n)/2$ (for $0\leqslant i\leqslant n-1$). This term (as a function of i) is maximized by $i=\pm n/4$. Therefore o is a vertex in $V^{(f)}_{\mathscr{A}}(S)$, shared by the n regions of pairs of points $(p_i, p_{(i+n/4)\bmod n})$. □

Now move from o to the right. First we enter the region of the points p_1 and p_2 at distance $\pi/4$ above and below the negative side of the x-axis. Next we enter the region that corresponds to the points immediately above (resp., below) p_1 (resp., p_2). When we reach the circle, we are already in the region of the points $\pi/3$ above and below the $-x$-axis. Eventually, we reach the region that corresponds to the top and bottom points (at $\pi/2$ above and below the y-axis). This path traverses $n/8$ cells. Since there are n possible directions, along which we traverse different cells, we have in total $n^2/8$ distinct cells.

The upper bound is obtained as follows. The Voronoi region in $V^{(f)}_{\mathscr{A}}(S)$ that corresponds to a pair of points $p,q\in S$ is the intersection of regions containing points *farther* from (p,q) than from (r,s) (according to $\mathscr{A}$), for all pairs of points $r,s\in S$. The "(p,q)-region" (with respect to (r,s)) is a double-wedge in the plane, where each wedge is on a different side of the line ℓ_{pq}. Since a wedge is convex, there are at most two convex cells that correspond to (p,q), each is the intersection of $\binom{n}{2}-1$ wedges

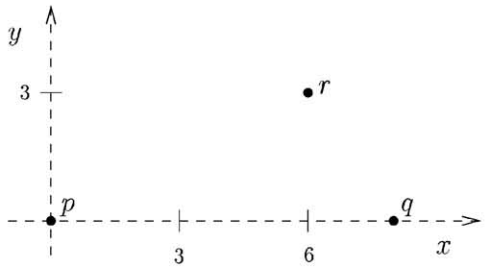

(a) Point set

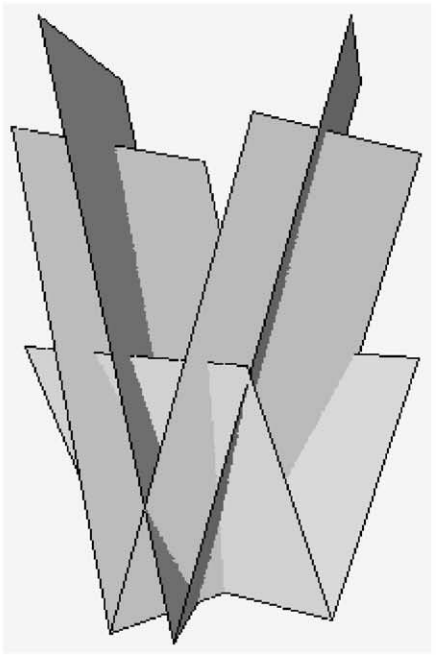

(b) Surfaces

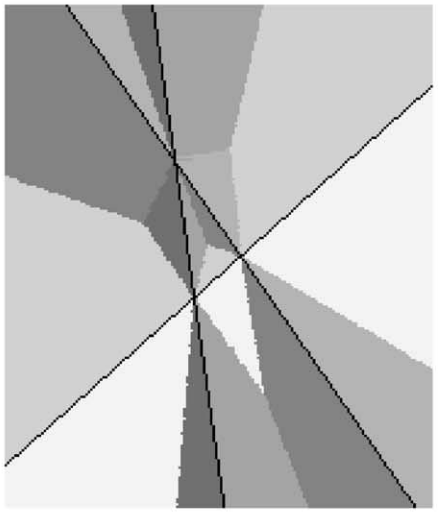

(c) Lower envelope

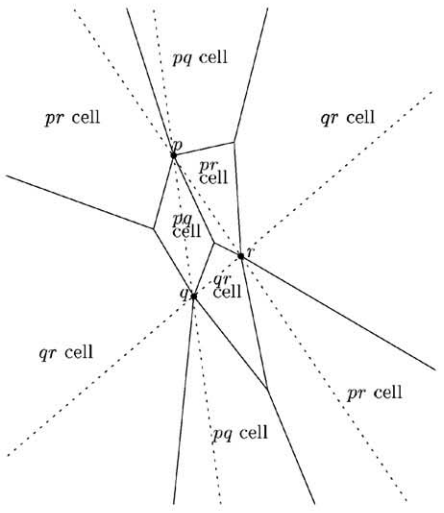

(d) Voronoi diagram

Fig. 3. $V^{(n)}_{\mathcal{A}}$ is the xy-projection of the lower envelope of a set of halfplanes.

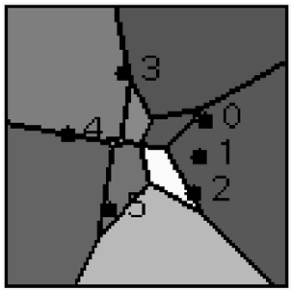

Fig. 4. Only hull points of S have regions in $V_{\mathscr{A}}^{(f)}(S)$.

on some side of ℓ_{pq}. In total the number of cells is at most twice the number of pairs of points, that is, $n(n-1)$.

Saying differently, $V_{\mathscr{A}}^{(f)}(S)$ is the upper envelope of $\binom{n}{2}$ planes in $\Re^3$, whose complexity is $\Theta(n^2)$ in the worst case [11, p. 216, Theorem 7.26]. □

We can further characterize the pairs of points that have nonempty regions in $V_{\mathscr{A}}^{(f)}(S)$ (see Fig. 4):

Theorem 7. (1) *Only pairs of points $p,q \in S$ where both p and q are vertices of $CH(S)$ (but not internal to an edge of the hull) have nonempty regions in $V_{\mathscr{A}}^{(f)}(S)$.*

(2) *Only pairs of vertices $p,q \in CH(S)$ that are antipodal to each other have infinite cells in $V_{\mathscr{A}}^{(f)}(S)$.*

Proof. (1) Assume that $V_{\mathscr{A}}^{(f)}(p,q) \neq \emptyset$ for $q \notin \mathrm{CH}(S)$. Let v be a point in $V_{\mathscr{A}}^{(f)}(p,q)$. Draw through q the line ℓ parallel to ℓ_{pv}. Denote by q' a vertex of $\mathrm{CH}(S)$ on the side of ℓ that does not contain p and v. It is easily seen that $\mathscr{A}(v,(p,q)) < \mathscr{A}(v,(p,q'))$, contradicting the assumption that $v \in V_{\mathscr{A}}^{(f)}(p,q)$.

A similar argument shows that a point site q on $\mathrm{CH}(S)$, which is not a vertex of the hull, cannot belong to a pair of points that has a nonempty region in $V_{\mathscr{A}}^{(f)}(S)$.

(2) Refer to Fig. 5. Let L_p and L_q be two parallel lines supporting $\mathrm{CH}(S)$ at the points p and q, respectively. Assume without loss of generality that ℓ_{pq} is horizontal and that $|\overline{pq}| = 1$. Let v be a point in the infinite strip bounded by L_p and L_q, and let h be the distance from v to ℓ_{pq}. Thus

$$\mathscr{A}(v,(p,q)) = h/2. \tag{1}$$

Consider a pair of points $r,s \in S$ out of which at least one point lies strictly inside the strip. (The other point may even be p or q.) Assume first that ℓ_{rs} is not vertical. Denote by θ the angle between ℓ_{pq} and ℓ_{rs}. Obviously, $|\overline{rs}| < 1/\cos\theta$. For ease of notation, set $\varepsilon = |\overline{rs}|\cos\theta$, where $0 < \varepsilon < 1$. Let h' be the distance from v to ℓ_{rs}. Thus

$$\mathscr{A}(v,(r,s)) = (h'\varepsilon)/(2\cos\theta). \tag{2}$$

Denote by L_v the line passing through v in parallel to L_p and L_q. Let h'' be the signed distance between the intersection of L_v and ℓ_{rs} to the line ℓ_{pq}. (In particular,

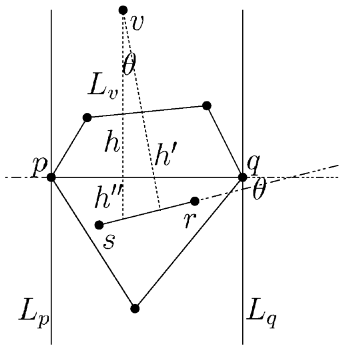

Fig. 5. Infinite cells in $V_{\mathscr{A}}^{(f)}(S)$ belong to antipodal points.

h'' is positive (resp., negative) if $L_v \cap \ell_{rs}$ lies below (resp., above) ℓ_{pq}.) The angle formed at v between L_v and the perpendicular from v to ℓ_{rs} is also θ. Therefore

$$\cos\theta = h'/(h + h''). \tag{3}$$

By combining Eqs. (2) and (3) we obtain

$$\mathscr{A}(v,(r,s)) = (h + h'')\varepsilon/2. \tag{4}$$

Note that by moving v away from ℓ_{pq}, the distance h'' remains unchanged, while we can make h arbitrarily large. We choose then h so that $h > h''/(\varepsilon(1-\varepsilon))$ for *all* possible choices of the points r and s. But then $(h+h'')\varepsilon/2 < h/2$, that is (by using Eqs. (1) and (4)), $\mathscr{A}(v,(r,s)) < \mathscr{A}(v,(p,q))$. This implies that v belongs to the Voronoi region of (p,q). (Or v belongs to the region of another pair of antipodal points whose supporting strip contains v. But when we move far enough along L_v away from ℓ_{pq}, the point v is eventually contained by only the strip bounded by L_p and L_q.)

The other case, where ℓ_{rs} is vertical, is straightforward. Here we always have $\mathscr{A}(v,(r,s)) \leqslant 0.5$. Therefore, we only need to set $h > 1$ to obtain $\mathscr{A}(v,(p,q)) > \mathscr{A}(v,(r,s))$. □

3.4. Computing $V_{\mathscr{A}}^{(n|f)}(S)$

We obtain $V_{\mathscr{A}}^{(n|f)}(S)$ by applying the general divide-and-conquer algorithm of [11, pp. 202–203] for computing the lower (or upper) envelope of a collection of bivariate functions. The merging step of this algorithm uses the standard line-sweep procedure of [10]. The total running time of the algorithm is $O((|M| + |M_1| + |M_2|)\log N)$, where M, M_1, and M_2 are the complexities of the envelopes and the two subenvelopes, respectively, and N is the number of surfaces. Since $M = O(n^4)$ and $N = O(n^2)$, we can compute $V_{\mathscr{A}}^{(n|f)}(S)$ in $O(n^4 \log n)$ time. The space required by the algorithm is dominated by the output size. Therefore the algorithm requires $O(n^4)$ space.

4. Distance from a line or a segment

We now define the 2-site distance to a line (or a segment) distance functions:

Definition 8. Given two points p,q in the plane, the "distance" $\mathscr{L}(v,(p,q))$ from a point v in the plane to the unordered pair (p,q) is defined as $\min_{u\in\ell_{pq}} d(v,u)$, the orthogonal distance from v to the line defined by p and q. Similarly, $\mathscr{G}(v,(p,q))$ is defined as $\min_{u\in\overline{pq}} d(v,u)$, the minimum distance from v to a point on the line-segment $\overline{pq}$.

The function $\mathscr{L}$ is very similar to the function $\mathscr{A}$, with the only difference that all the strips around all pairs of point sites grow at the same speed, irrespective of the distance between the two points of each pair. Hence, all the Voronoi surfaces are pairs of halfplanes bounded by lines on the plane $z=0$ and ascending outward of it with a slope of 45°. The analyses of the complexities of their lower and upper envelopes are identical to those of the triangle-area distance function (see Section 3). Therefore, given a set S of n points, the complexities of $V^{(n)}_{\mathscr{L}}(S)$ and $V^{(f)}_{\mathscr{L}}(S)$ are $\Theta(n^4)$ and $\Theta(n^2)$, respectively.

We turn our attention, then, to $V^{(n|f)}_{\mathscr{G}}(S)$. The growing shape that corresponds to the function $\mathscr{G}$ is a hippodrome, a rectangular shape centered about the line segment connecting two point sites, and expanded by two hemicycles, attached to the far ends of the shape, with diameter equal to the width of the rectangle.

4.1. Nearest-neighbor diagram

Theorem 9. *The combinatorial complexity of* $V^{(n)}_{\mathscr{G}}(S)$ *is* (*for all sets of points in general position*) $\Theta(n^4)$.

Proof. The n points of S define $\Theta(n^2)$ segments which always have $\Theta(n^4)$ intersection points. This is a consequence of the fact that every planar drawing of a graph with n vertices and $m \geqslant 4n$ edges (without self or parallel edges) has $\Omega(m^3/n^2)$ crossing points [1,8]. (In our case $m=\binom{n}{2}$.) All these intersection points are features of $V^{(n)}_{\mathscr{G}}(S)$. Hence the lower bound.

The upper bound is obtained by splitting each segment at each intersection point with another segment. The complexity of $V^{(n)}_{\mathscr{G}}(S)$ is upper bounded by the complexity of the nearest-neighbor Voronoi diagram of the set of "broken" segments. Since the latter set consists of $\Theta(n^4)$ nonintersecting segments (except in their endpoints), its complexity is $\Theta(n^4)$ [7]. □

The diagram $V^{(n)}_{\mathscr{G}}(S)$ can be computed in $O(n^4 \log n)$ time and $O(n^4)$ space by the lower-envelope algorithm of [11, pp. 202–203] or by the special-purpose algorithms of Fortune [4] and Yap [13] with the same asymptotic running time and space complexities.

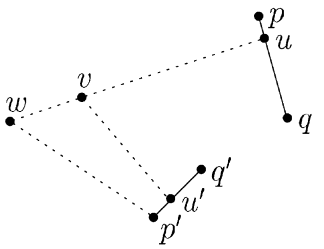

Fig. 6. Cells in $V_{\mathscr{G}}^{(f)}(S)$ are infinite.

4.2. Furthest-neighbor diagram

We need two lemmas for analyzing the structure of $V_{\mathscr{G}}^{(f)}(S)$:

Lemma 10. *If the region of two point sites p, q in $V_{\mathscr{G}}^{(f)}(S)$ is nonempty, then p and q are the two extreme points on one side of some direction.*

Proof. Let the region of $p,q \in S$ be nonempty, so that there exists some point $v \in \Re^2$ such that $v \in V_{\mathscr{G}}^{(f)}(p,q)$. Assume without loss of generality that p is closer to v than q, and let u be the point on $\overrightarrow{pq}$ closest to v. We claim that p and q are the furthest sites along the direction $\overrightarrow{vu}$. Assume to the contrary that there exists a site r further than p (or from p and q, if $u \notin \{p,q\}$) along $\overrightarrow{vu}$. But in this case v must belong to the region of (q,r) in $V_{\mathscr{G}}^{(f)}(S)$, contradicting the assumption that $v \in V_{\mathscr{G}}^{(f)}(p,q)$. □

Lemma 11. *Every cell in $V_{\mathscr{G}}^{(f)}(S)$ is infinite.*

Proof. As in the proof of Lemma 10, refer to the region of (p,q) in $V_{\mathscr{G}}^{(f)}(S)$. Let v be a point in this region, and let u be the point on $\overline{pq}$ closest to v, so that

$$\mathscr{S}(v,(p,q)) = d(v,u). \tag{5}$$

We show that the ray originating at v and pointing *away* from u (see Fig. 6) is fully contained in $V_{\mathscr{G}}^{(f)}(p,q)$. Assume to the contrary that there exists a point w on that ray, which belongs to the region of (p',q'), so that

$$\mathscr{S}(w,(p,q)) < \mathscr{S}(w,(p',q')). \tag{6}$$

Denote by u' the point on $\overline{p'q'}$ closest to v, so that

$$\mathscr{S}(v,(p',q')) = d(v,u'). \tag{7}$$

We have

$$\mathscr{S}(w,(p,q)) = d(w,v) + d(v,u). \tag{8}$$

By substituting (5) in (8) we obtain:

$$\mathscr{S}(w,(p,q)) = d(w,v) + \mathscr{S}(v,(p,q)). \tag{9}$$

From the Euclidean triangle inequality (and in view of (7)) we have:

$$\mathscr{S}(w,(p',q')) \leqslant d(w,u') < d(w,v) + d(v,u') = d(w,v) + \mathscr{S}(v,(p',q')). \qquad (10)$$

Finally, by rewriting (9) and using (6) and (10), we find that

$$\mathscr{S}(v,(p,q)) < \mathscr{S}(v,(p',q')),$$

which is a contradiction to the assumption that $v \in V_{\mathscr{G}}^{(f)}(p,q)$. Hence also $w \in V_{\mathscr{G}}^{(f)}(p,q)$. □

Theorem 12. *The combinatorial complexity of $V_{\mathscr{G}}^{(f)}(S)$ is $\Theta(n)$ in the worst case.*

Proof. Lemmas 10 and 11 dictate the structure of $V_{\mathscr{G}}^{(f)}(S)$. Let $p_1, p_2, \ldots, p_{h_1}$ be the sequence of points (say, clockwise) of S along $\mathrm{CH}(S)$, where h_1 is the number of (so-called "outer") hull points of S. Let $q_1, q_2, \ldots, q_{h_2}$ (so-called "inner" hull points) be the vertices of the convex hull of the remaining set, that is, of $\mathrm{CH}(S \setminus \mathrm{CH}(S))$. Lemma 10 allows two types of nonempty regions in $V_{\mathscr{G}}^{(f)}(S)$:

(1) Regions corresponding to pairs of consecutive points along $\mathrm{CH}(S)$, that is, of (p_1, p_2), $(p_2, p_3), \ldots, (p_{h_1}, p_1)$.
(2) Regions of pairs of the form (p_i, q_j), where $1 \leqslant i \leqslant h_1$ and $1 \leqslant j \leqslant h_2$. Due to Lemma 10 it is mandatory that for a fixed outer-hull point p_i, all the inner-hull points q_j, for which the region of (p_i, q_j) in $V_{\mathscr{G}}^{(f)}(S)$ is nonempty, belong to some continuous range of points along the inner hull of S, say, $q_{j'}, q_{j'+1}, \ldots, q_{j''}$ (where indices are taken modulo h_2). Moreover, only $q_{j'}$ (resp., $q_{j''}$) may have with p_{i-1} (resp., p_{i+1}) a nonempty region in $V_{\mathscr{G}}^{(f)}(S)$, for otherwise the two-extreme-points property would be violated.

The structure of $V_{\mathscr{G}}^{(f)}(S)$ is now obvious: traversing the diagram rotationally around $\mathrm{CH}(S)$ (far enough from $\mathrm{CH}(S)$ so as to pass through all the cells), we alternate between regions of the first type (described above) to (possible) ranges of regions of the second type. Namely, we are guaranteed to go through nonempty regions of (p_{i-1}, p_i) and (p_i, p_{i+1}) (where indices are taken modulo h_1), possibly separated by a range of regions of $(p_i, q_{j'})$, $(p_i, q_{j'+1}), \ldots, (p_i, q_{j''})$ (where indices of the inner-hull points are taken modulo h_2). As noted above, $q_{j''}$ is the only point in the range $j' \leqslant j \leqslant j''$ that may have, together with p_{i+1}, a nonempty region in $V_{\mathscr{G}}^{(f)}(S)$.

This structure is shown in Fig. 7. Points 0–2 belong to the outer hull of the point set, whereas points 3–6 belong to the inner hull. The clockwise order of $V_{\mathscr{G}}^{(f)}$ contains the cells of $(0,1)$ and $(1,2)$, separated by the cells of $(1,3)$, $(1,4)$, $(1,5)$, and $(1,6)$. The cell of $(0,3)$ appears before that of $(0,1)$, and the cell of $(2,6)$ appears after that of $(1,2)$.

It is now easy to upper bound the number of cells in $V_{\mathscr{G}}^{(f)}(S)$. There are exactly h_1 regions of the first type and at most $h_1 + h_2$ regions of the second type. Since $1 \leqslant h_1, h_2 \leqslant n$ and $h_1 + h_2 \leqslant n$, the diagram $V_{\mathscr{G}}^{(f)}(S)$ contains at most $2n$ nonempty regions. In the worst case, then, the complexity of $V_{\mathscr{G}}^{(f)}(S)$ is $\Theta(n)$. However, the above discussion also implies trivial constructions of arbitrary number of points, for which the complexity of $V_{\mathscr{G}}^{(f)}(S)$ is constant. □

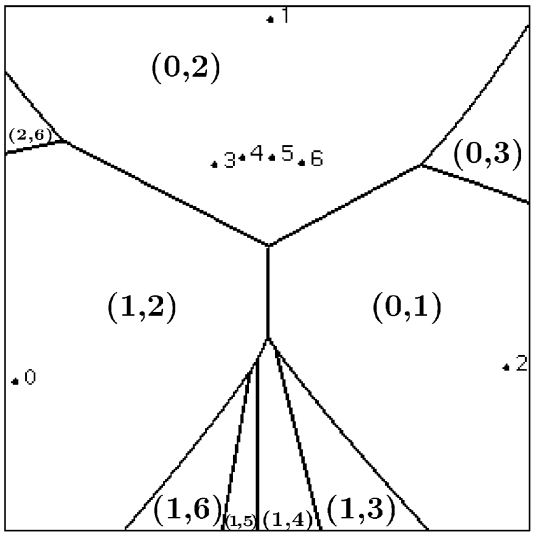

Fig. 7. $V_{\mathscr{D}}^{(f)}$ of seven points.

5. Difference between distances

We now define the 2-site difference-between-distances distance function:

Definition 13. Given two points p,q in the plane, the "distance" $\mathscr{D}(v,(p,q))$ from a point v in the plane to the unordered pair (p,q) is defined as $|d(v,p)-d(v,q)|$.

For a fixed pair of points p and q, the curve $\mathscr{D}(v,(p,q))=c$, for a constant $c \geqslant 0$, is a pair of quadratic curves in the plane. For $c=0$ this is the bisector of the line segment $\overline{pq}$. As c grows, the "strip" widens but in a shape different than that corresponding to the area function. The borders of the strip advance as a pair of hyperbolas with the invariant that the points p and q always remain the foci of the hyperbolas. These borders form "beaks" closing away of the initial bisector of $\overline{pq}$. The function $\mathscr{D}(v,(p,q))$ reaches its maximum value $d(p,q)$ when the two beaks close on the two rays, one emanating from p away from q, and the other emanating from q away from p. (This is obtained from the triangle inequality.) The bisectors between pairs of regions in the diagram are quartic functions. Figs. 8(a) and (b) show two configurations of four points. The region of points closer to $(0,1)$ (according to $\mathscr{D}$) is shown in dark gray, while the region of $(2,3)$ is shown in light gray.

5.1. Nearest-neighbor diagram

As with the area distance function, since the $\Theta(n^4)$ intersection points of the $\Theta(n^2)$ bisectors of the line segments joining pairs of points in S are all features in $V_{\mathscr{D}}^{(n)}(S)$,

Fig. 8. Bisectors for $\mathscr{D}$.

the combinatorial complexity of the diagram is $\Omega(n^4)$. We show a nearly-matching upper bound (but conjecture, however, that the correct bound is $\Theta(n^4)$).

Theorem 14. *The combinatorial complexity of* $V_{\mathscr{D}}^{(n)}(S)$ *is* $O(n^{4+\varepsilon})$ *(for any* $\varepsilon > 0$*) in the worst case.*

Proof. The collection of $\Theta(n^2)$ surfaces $F_{\mathscr{D}}^{p,q}$ fulfills Assumptions 7.1 of [11, p. 188]:

(i) Each surface is an algebraic surface of maximum constant degree.
(ii) Each surface is totally defined (this is stronger than needed);
(iii) Each triple of surfaces intersect in at most a constant number of points. (This follows from Bézout's theorem.)

Hence we may apply Theorem 7.7 of [11, p. 191] and obtain the claimed complexity of $V_{\mathscr{D}}^{(n)}(S)$.

As with the area distance function, we apply the same divide-and-conquer algorithm with a plane-sweep for the merging step (Theorem 7.16 of [11, p. 203]). In this case $M = O(n^{4+\varepsilon})$ and $N = O(n^2)$. Thus we obtain an $O(n^{4+\varepsilon} \log n)$-time and $O(n^{4+\varepsilon})$-space algorithm for computing $V_{\mathscr{D}}^{(n)}(S)$. □

5.2. Furthest-neighbor diagram

Theorem 15. *The combinatorial complexity of* $V_{\mathscr{D}}^{(f)}(S)$ *is* $\Theta(n^2)$ *in the worst case.*

Proof. For every point $v \in \Re^2$, the pair $p, q \in S$ for which $\mathscr{D}(v,(p,q))$ is maximized must consist of the nearest and the furthest neighbors of v according to the regular Euclidean distance function. Hence, $V_{\mathscr{D}}^{(f)}(S)$ is the overlay of the regular nearest- and furthest-neighbor Voronoi diagrams of S. The complexity of this overlay is $\Theta(n^2)$ in the worst case. □

The overlay of the nearest- and furthest-neighbor Voronoi diagrams can be computed in $O(n^2)$ time [2], or in an output-sensitive manner in $O(n \log n + k)$ time, where k is the complexity of the overlay, by the algorithm of Guibas and Seidel [5].

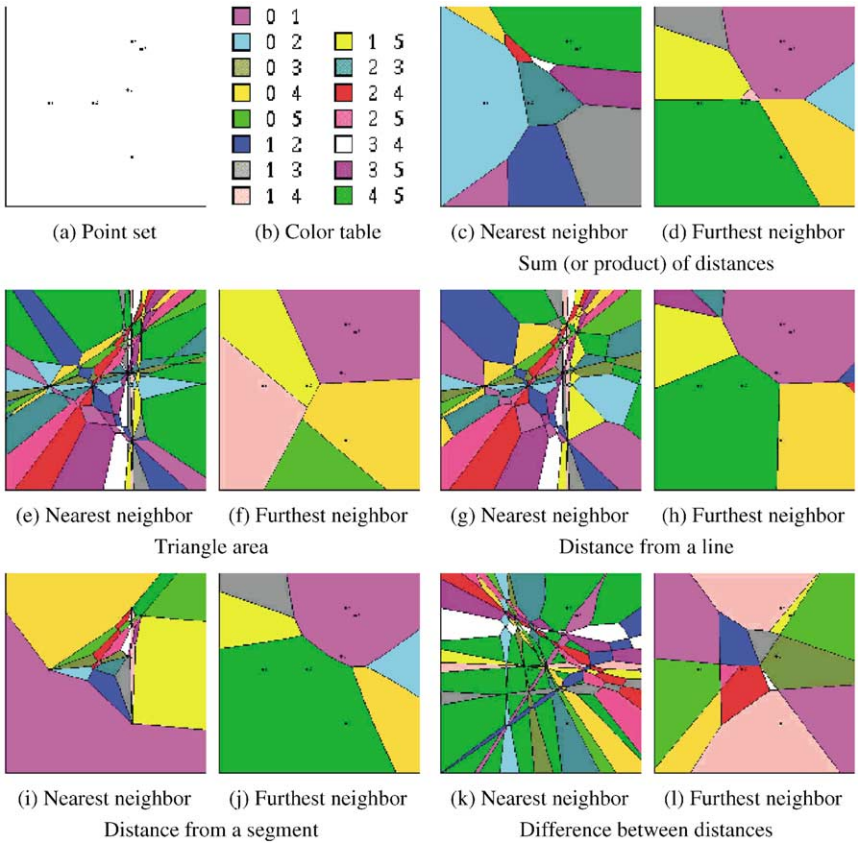

Fig. 9. 2-site Voronoi diagrams of a set of six points.

6. Conclusion

In this paper, we present a new notion of distance between a point and a pair of points, which, to the best of our knowledge, was never discussed in the literature, and define a few instances of it. For each such distance function we investigate the nearest- and furthest-neighbor Voronoi diagrams of a set of points in the plane and methods for computing it.

We have implemented a Java applet that computes the Voronoi diagrams with respect to the distance functions discussed in this paper (and some more functions), and a Web page which provides interface to this applet. The applet supports interactive selection of the point set and on-line computation and display of the Voronoi diagrams. Fig. 9(a) shows a set of six points. Fig. 9(b) shows the color table used in the following Voronoi diagrams. Figs. 9(c–l) show the nearest- and furthest-neighbor diagrams of the six points for the sum (or product) of distances, triangle area, distance from a line, distance from a segment, and difference-between-distances distance functions, respectively.

Future research directions include:

(1) Investigating the respective diagrams of more distance functions, e.g., the perimeter of the triangle defined by the three points, the radius of the circle defined by them, etc.
(2) Generalizing 2-site distance functions to higher dimensions.
(3) Defining and characterizing distance functions from a point to more than 2 points.

Acknowledgements

We thank Micha Sharir for commenting on the use of lower envelopes for computing the diagrams discussed in this paper. The paper also benefited from discussions with several people at the second CGC Workshop on Computational Geometry (held in Durham, NC, in October 1997), in which it was first presented. In particular, the use of the zone theorem to improve the upper bound on the complexity of $V_{\mathscr{A}}^{(n)}(S)$ was suggested by Jeff Erickson. David Eppstein also made several useful suggestions during the workshop. We also thank Daniel Scharstein for suggesting the basic idea for proving Theorem 7.1.

References

[1] M. Ajtai, V. Chvátal, M. Newborn, E. Szemerédi, Crossing-free subgraphs, Ann. Discrete Math. 12 (1982) 9–12.
[2] H. Ebara, N. Fukuyama, H. Nakano, Y. Nakanishi, Roundness algorithms using the Voronoi diagrams, Proceedings of the First Canadian Conference on Computational Geometry, Vol. 41, 1989.
[3] H. Edelsbrunner, R. Seidel, M. Sharir, On the zone theorem for hyperplane arrangement, SIAM J. Comput. 2 (1993) 418–429.
[4] S. Fortune, A sweepline algorithm for Voronoi diagrams, Proceedings of the Second Annual ACM Symposium on Computational Geometry, 1986, pp. 313–322.
[5] L.J. Guibas, R. Seidel, Computing convolutions by reciprocal search, Discrete Comput. Geom. 2 (1987) 175–193.
[6] D.T. Lee, On k-nearest neighbor Voronoi diagrams in the plane, IEEE Trans. Comput. 31 (1982) 478–487.
[7] D.T. Lee, R.L. Drysdale, III, Generalization of Voronoi diagrams in the plane, SIAM J. Comput. 10 (1981) 73–87.
[8] F.T. Leighton, Complexity Issues in VLSI, MIT Press, Cambridge, MA, 1983.
[9] A. Okabe, B. Boots, K. Sugihara, Spatial Tessellations: Concepts and Applications of Voronoi Diagrams, Wiley, New York, 1992.
[10] F.P. Preparata, M.I. Shamos, Computational Geometry: An Introduction, Springer, Berlin, 1985.
[11] M. Sharir, P.K. Agarwal, Davenport-Schinzel Sequences and Their Geometric Applications, Cambridge University Press, Cambridge, 1995.
[12] M.R. Spiegel, Mathematical Handbook of Formulas and Tables, Schaum's Outline Series, McGraw-Hill, New York, 1990.
[13] C.-K. Yap, An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments, Discrete Comput. Geom. 2 (1987) 365–393.