

Voronoi Diagrams for Convex Polygon-Offset Distance Functions*

G. Barequet,¹ M. T. Dickerson,² and M. T. Goodrich¹

¹ Center for Geometric Computing,
Department of Computer Science, Johns Hopkins University,
Baltimore, MD 21218, USA
{barequet, goodrich}@cs.jhu.edu

² Department of Mathematics and Computer Science, Middlebury College,
Middlebury, VT 05753, USA
dickerso@middlebury.edu

Abstract. In this paper we develop the concept of a convex *polygon-offset* distance function. Using offset as a notion of distance, we show how to compute the corresponding nearest- and furthest-site Voronoi diagrams of point sites in the plane. We provide near-optimal deterministic $O(n(\log n + \log^2 m) + m)$ -time algorithms, where n is the number of points and m is the complexity of the underlying polygon, for computing compact representations of both diagrams.

1. Introduction

The Voronoi diagram of a set S of n points (called sites) in the plane 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. It is a well-known, powerful tool for handling a host of geometric problems dealing with distance relationships (see, e.g., [OBS]). Voronoi diagrams have been used extensively, for example, for solving nearest-neighbor, furthest-neighbor, and matching problems in many contexts. The underlying distance function used to define Voronoi diagrams is typically either the usual Euclidean metric, or more generally a distance function based upon one of the L_p metrics. There has also been some interesting work done using convex (scaling) distance functions (also

* Work on this paper by the first and the third authors has been supported in part by the U.S. Army Research Office under Grant DAAH04-96-1-0013. Work by the second author has been supported in part by the National Science Foundation under Grants CCR-93-1714 and CCR-99-02032. Work by the third author has been supported also by NSF grant CCR-96-25289.

called Minkowski functionals [KN, p. 15]), which are extensions of the scaling notion for circles (in the Euclidean case) to convex polygons.

Voronoi diagrams also play an important role in tolerancing problems that arise in manufacturing processes of physical three-dimensional objects. However, defining distance in terms of an offset from a polygon¹ is actually more natural than scaling in this application, as well as in many others. This is because the local error of a production tool (a milling head, a laser beam, etc.) is independent of the location of the produced feature relative to some artificial reference point (the origin). The production error at any location on the object's boundary is not likely to depend on a scale factor from some arbitrary center point, such as the scaling center. Instead, tolerances are more likely to be given as an allowable error distance from the boundary. Thus, offset polygons and not scaled polygons are the correct way to model manufacturing tolerances [Cr]. In a companion paper [BBDG] we solve a few geometric problems that capture the notion of tolerancing. The algorithms proposed there make use of the Voronoi diagram. This motivates the computation of the Voronoi diagram of a point set in the plane, whose underlying distance function is based on offsetting.

In this paper we investigate distance functions based on *offsetting* convex polygons. While the scaling operation shifts each edge of the polygon proportionally to its distance from the origin, the offset operation shifts all the edges by the same amount. Offset polygons are therefore not homothetic copies of the original polygon (unless the original polygon is regular). We are interested in the investigation of basic properties of polygon-offset distance functions, with particular attention paid to how they may be used in the definition and computation of Voronoi diagrams. This is not obvious, since the polygon-offset functions are not metrics.

1.1. Related Previous Work

We are not familiar with any prior work on defining distance in terms of offset polygons, nor in methods for defining Voronoi diagrams in terms of such functions. Minkowski (see, e.g., [KN]) was the first to study the related notion of what are now commonly referred to as *convex distance functions*. He showed, for example, that distance can be consistently defined in terms of a scaling of a convex polygon containing the origin, and that while such functions do not in general define metrics (since they are usually not symmetric), they exactly characterize the distance functions satisfying the triangle inequality. Chew and Drysdale [CD] brought this notion into the computational geometry literature, showing that one can consistently define nearest-neighbor Voronoi diagrams using convex distance functions. They give an $O(nm \log n)$ -time method for constructing such diagrams for a set S of n points in the plane, with distance defined by a scaling of an m -edge convex polygon. This is actually quite close to optimal for fully constructing the Voronoi diagram, as they show that such a diagram can be of size $\Theta(nm)$. Klein, Mehlhorn, Meiser, and others [Kl], [KMM], [KW], [MMO], [MMR] generalized this work even further, showing how to define Voronoi diagrams in a very abstract setting. They also give

¹ Intuitively speaking, offsetting a polygon is done by locally shifting all its edge. We give a precise definition later in the paper.

randomized incremental constructions for this abstract setting that can be applied to nearest- and furthest-neighbor Voronoi diagrams for convex distance functions. Strictly speaking, the running times of these constructions are expected to be $O(nm \log n)$, but it is possible to use their approaches to construct “compact” Voronoi diagram representations in $O(n \log n \log m + m)$ expected time. For the case of nearest-neighbor diagrams this was improved by McAllister et al. [MKS], who give a deterministic $O(n(\log n + \log m) + m)$ -time method for constructing a compact Voronoi diagram complete enough for answering nearest-neighbor queries for a given convex distance function in $O(\log n + \log m)$ time. They do not address furthest-neighbor diagrams, however, nor do they address the case when distance is defined by offsets of a convex polygon.

Aichholzer and Aurenhammer [AA] present an algorithm for constructing the so-called *straight skeleton* of a collection of polygonal chains, which identifies with the medial axis of a convex polygon. They mention the use of the skeleton as a distance function and give a sketch [ibid., Fig. 4] of a Voronoi diagram of two straight segments. They take the set of polygonal chains that define the distance function to be also the set of sites for which the Voronoi diagram is sought. We, however, make the distinction between the convex polygon that defines the distance function and the set of point sites in the diagram. Aichholzer et al. [AAAG] further study the properties and applications of the straight skeleton of a simple polygon. They remark that the skeleton of a polygon “is no Voronoi-diagram-like structure.” Unlike in the cited work, the constructions presented in this paper are Voronoi diagrams of *point sets* and not of polygons; the polygon only defines the underlying distance function. We are therefore concerned with a very different set of problems and constructions.

1.2. Our Results

In this paper we formally develop the concept of a convex polygon-offset distance function and explore its properties. One such interesting property is the fact that polygon-offset distance functions do *not* in general satisfy the triangle inequality. This, of course, follows from the contra-positive of Minkowski’s characterization theorem, but we provide a simple constructive proof. In spite of this fact, however, we show that convex polygon-offset distance functions nevertheless satisfy all the topological properties for abstract Voronoi diagrams [KI], [KMM], [KW], [MMO], [MMR]. Finally, given a set S of n points in the plane, we show how to construct compact representations of nearest- and furthest-site Voronoi diagrams deterministically for S with respect to an offset distance defined by an m -edge convex polygon in $O(n(\log n + \log^2 m) + m)$ time. We use the tentative prune-and-search paradigm of Kirkpatrick and Snoeyink [KS] to design some geometric primitives. While McAllister et al. [MKS] designed similar primitives with respect to the Euclidean distance functions, we do that with respect to the convex polygon-offset distance function. These primitives are used in generalizations of Fortune’s and Rappaport’s algorithms for computing the nearest- and furthest-neighbor Voronoi diagrams, respectively.

The paper is organized as follows. In Section 2 we give some notations and definitions. We describe in Section 3 how to construct the nearest-site Voronoi diagram, and show in Section 4 how to obtain its compact representation. Analogously, we describe in Section 5

how to compute the furthest-site Voronoi diagram in full and in compact representations. We terminate in Section 6 with some concluding remarks.

2. Preliminaries

The usual Euclidean distance between two points $p, q \in \mathbb{R}^2$ can be modeled as follows. Place the unit circle centered at p and scale it until q occurs at the boundary of the scaled circle. The radius of the latter circle is the distance between p and q . This distance function is obviously symmetric. In this section we define the convex polygon-offset distance function and show that in general it does not fulfill the triangle inequality.

2.1. Metrics and Convex Distance Functions

A metric on the plane (see [KW, p. 16] and [Du]) is a mapping $D: \mathbb{R}^2 \times \mathbb{R}^2 \rightarrow \mathbb{R}_{\geq 0}$ such that for any points $p, q, r \in \mathbb{R}^2$ the following conditions are fulfilled:

1. $D(p, q) = 0$ if and only if $p = q$ (identity),
2. $D(p, q) = D(q, p)$ (symmetry), and
3. $D(p, r) \leq D(p, q) + D(q, r)$ (triangle inequality).

For every convex polygon P (which contains the origin) we can define the convex distance function D_P as follows. Let p and q be points in the plane as before. Translate P by p (that is, place the origin at p) and scale it so that q occurs at the boundary of P . The distance $D_P(p, q)$ is the respective scaling factor of P . The function D_P is not symmetric (thus it is not a metric) unless P is symmetric about the origin [CD]. Convex distance functions are normalized such that:

1. $D_P(p, p) = 0$ for every point p and polygon P .
2. $D_P(p, q) = 1$ if and only if $q \in \partial\tau_p(P)$, where τ_p is the translation of the plane that maps the origin to p .

Equivalently, we can define D_P in terms of the normal Euclidean distance function E : $D_P(p, q) = E(p, q)/E(p, q')$, where q' is the intersection of the vector \vec{pq} with the (unit) polygon P . Note that by this definition, different translations of a convex polygon P (as long as they all contain the origin) define different convex distance functions. However, this is not the case with the polygon-offset distance function.

2.2. Convex Polygon-Offset Distance Functions

We first define the *offset* of a convex polygon. A convex polygon P is the intersection of a collection of closed halfplanes $\{H_i\}$, each defined by an edge of P . The offset polygon $O_{P,\varepsilon}$ is the intersection of $\{H_i(\varepsilon)\}$, where $H_i(\varepsilon)$ is the halfplane parallel to H_i with bounding line translated by distance ε . Positive (resp., negative) values of ε stand for translating the edges outward (resp., inward) from the polygon. Figure 1 shows an inner and an outer offset of a convex polygon.




Fig. 1. Inner and outer ϵ -offsets of a convex polygon.

None of the halfplanes $\{H_i\}$ is redundant in the definition of P . This is also true for $O_{P, \epsilon}$ where $\epsilon > 0$. However, for negative values of ϵ , some or all of $\{H_i(\epsilon)\}$ may become redundant. The value $\epsilon_0 < 0$ for which O_{P, ϵ_0} degenerates into a point is the radius of the largest circle inscribed into P . We call this point the offset center of P .

There is a strong relation between the continuous change of $O_{P, \epsilon}$ (as a function of ϵ) and the medial axis of P . Recall that the polygon P is convex. The medial axis of a polygon P is the set of points inside P that have more than one closest point among the points of ∂P [Or, p. 193]. As a direct consequence of this definition, we have that changing ϵ continuously makes the vertices of P slide along the edges of the medial axis. Assume we start with some $\epsilon > 0$ and reduce it continuously, thus we slide the vertices of P inward along its medial axis. When two neighboring vertices $v_i, v_j \in P$ meet at a vertex of the medial axis, the halfplane that corresponds to the edge of P that connects between v_i and v_j becomes redundant, and v_i, v_j are merged (on the offset polygon) into one vertex. At some value of $\epsilon > \epsilon_0$ we are left with a triangle (or with a regular k -gon, for some $k \geq 3$), which, when ϵ reaches ϵ_0 , is offset to a point, the center of P . A degenerate case occurs when P contains two parallel edges that define the width of P . In this case O_{P, ϵ_0} is a line segment rather than a point. For our purposes we can choose any arbitrary point in that segment, say, its center point, to be the center of P . Figure 2 illustrates the relation between offset of P and its medial axis.

We are now ready to define the convex polygon-offset distance function D_P . Let p and q be two points in the plane. Consider the value of ϵ for which a translate of $O_{P, \epsilon}$




Fig. 2. Offset of a convex polygon based on its medial axis. (a) Point center and (b) segment center.




Fig. 3. Dominating edges for the convex polygon-offset distance function.

centered at p contains q on its boundary. The offset distance is defined as

$$D_P(p, q) = (\varepsilon + |\varepsilon_0|)/|\varepsilon_0| = \varepsilon/|\varepsilon_0| + 1.$$

This definition normalizes $\mathcal{D}_P(p, p)$ so that $\mathcal{D}_P(p, p) = 0$ and $\mathcal{D}_P(p, q) = 1$ for every point p in the plane and point q on the boundary of P centered at p .

2.3. Basic Properties of \mathcal{D}_P

As defined in the previous section, for measuring the P -offset distance from a point p to another point q , we place P such that its offset center coincides with p . We then offset P (inward or outward) until it hits q . We say that the edge of P that hits q *dominates* it. Consider the subdivision of the plane induced by extending the edges of the medial axis of P (that are coincident to vertices of P) to infinity. Each region in this subdivision is dominated by an edge of P ; furthermore, there is a 1-to-1 mapping between regions of the subdivision and edges of P . When we move away from p along the ray \vec{pq} we may cross an edge of the medial axis. This event corresponds to switching from the dominance of some polygon edge to that of another polygon edge. Figure 3 illustrates this situation: the ray \vec{pq} is first dominated by $e_1 \in P$ (in the inner offsets of P , close enough to p , the offset halfplane that corresponds to the edge e_2 is redundant), and then it crosses an edge of the medial axis of P and becomes dominated by e_2 .

We now give a more specific relationship between \mathcal{D}_P and the Euclidean distance function. Refer to Fig. 4. Let ℓ_e be the direction in which the dominating edge e is offset outward of P , and let θ be the angle between ℓ_e and \vec{pq} . A unit offset of e (that is, translating e by 1 along ℓ_e) corresponds to a $(1/\cos \theta)$ -move along \vec{pq} .² Our goal is to measure distance along \vec{pq} in terms of the corresponding offset of e . When $\theta = 0$, we have $1/\cos \theta = 1$. As θ increases, a unit offset of e results in an increasing Euclidean distance along \vec{pq} , and as θ approaches $\pi/2$, the distance along \vec{pq} that corresponds to a unit offset of e grows to infinity. This is because an infinitesimal offset of e is then sufficient for any arbitrarily large move along \vec{pq} .

It is clear that within each region of the plane subdivision induced by the medial axis, the quantity $(1/\cos \theta)$ is constant in any direction (that is, for any θ). Thus, the

² Note the difference between a unit offset and a unit of \mathcal{D}_P . The latter measure is normalized, so that a unit of \mathcal{D}_P corresponds to the distance between the center of P and any point on ∂P .




Fig. 4. The relation between \mathcal{D}_P and the Euclidean distance.

offset distance is linear in the Euclidean distance (with the multiplicative factor $\cos \theta$). It turns out that the quantity $(1/\cos \theta)$ along any direction can only decrease. This quantity changes only when the ray \vec{pq} crosses an edge of the medial axis. That is, it exits a region that corresponds to some edge $e_1 \in P$ and enters a region that corresponds to another edge $e_2 \in P$. The edge of the medial axis is the bisector of e_1 and e_2 . Elementary geometry shows that \vec{pq} forms a smaller angle with the normal to e_2 than the angle with the normal to e_1 . Thus we have:

Theorem 1. *If p, s, q are collinear points in this order, then $\mathcal{D}_P(p, s) + \mathcal{D}_P(s, q) \leq \mathcal{D}_P(p, q)$ and equality holds if and only if the entire segment $[pq]$ is dominated by the same edge of P when P is centered at p .*

Proof. Assume that there is a change of the dominating edge along the segment $[pq]$ and denote the point of transition by r . Since the quantity $(1/\cos \theta)$ of \vec{pq} can only decrease at r (i.e., the offset distance with respect to \mathcal{D}_P can only increase) by such a change, it is imperative that the offset distance (per a unit Euclidean distance) along $[pr]$ is smaller than that along $[rq]$. Now choose any point $s \in [rq]$. When measuring the distance along $[sq]$ we are first charged according to the cheaper rate (since now P is centered at s). Therefore $\mathcal{D}_P(p, s) + \mathcal{D}_P(s, q) < \mathcal{D}_P(p, q)$. A similar argument holds for the case $s \in [pr]$. Equality holds when there is no such point r , that is, there is no change in the dominating edge along $[pq]$. \square

Corollary 2. *Convex polygon-offset distance functions defined by nonregular polygons do not fulfill the triangle inequality.*

Note that Minkowski [KN, p. 15, Theorem 2.3] proved that given a distance function based on scaling a shape S , the triangle inequality holds if and only if S is convex.

3. Nearest-Site Voronoi Diagram

We now address the construction of a nearest-site Voronoi diagram using the polygon-offset distance function \mathcal{D}_P (for some convex polygon P). Given a planar point set S , we denote its nearest- and furthest-site Voronoi diagrams with respect to \mathcal{D}_P by $\mathcal{V}_P^N(S)$ and $\mathcal{V}_P^F(S)$, respectively. When the set S is clear from the context, we omit it from this notation. Figure 5(a) shows a convex polygon P , while the Voronoi diagram of three points with respect to \mathcal{D}_P is shown in Fig. 5(b).




Fig. 5. A nearest-site Voronoi diagram. (a) Underlying polygon and (b) diagram of three points.

It is well known that for any scaled-polygon distance function, the Voronoi cell of a point p is always a starshaped polygon whose kernel contains p . In contrast, there exist nearest-neighbor Voronoi cells for the polygon-offset distance function which are nonstarshaped. Figure 6 shows a quadrangle P (with dashed edges) which defines a convex-offset distance function. Its medial axis is shown with dotted lines. The figure also shows the respective Voronoi diagram of three points p, q, r equally spaced along a line. The Voronoi cell that corresponds to the site q is not starshaped.

For computing the Voronoi diagram, we begin by showing that this notion of distance fits the unifying approach of Klein for abstract Voronoi diagrams [KW], [K1]. Klein




Fig. 6. A nonstarshaped cell of q .

proposes to replace the distance notion by *bisecting curves*. Each pair of sites, regardless of site shapes or distance function, are separated by curves which divide the plane into two portions, each corresponding to one site. The Voronoi region of a site s is the intersection of all the s -portions defined by bisecting curves between s and all the other sites.

3.1. More Properties of \mathcal{D}_P

In order to show that the convex polygon-offset distance function is valid for defining Voronoi diagrams, we need to define precisely the nearest-site Voronoi diagram in this context: the Voronoi cell that corresponds to a site p consists of all the points x in the plane for which $\mathcal{D}_P(x, p) \leq \mathcal{D}_P(x, q)$ for all $q \neq p$. Note that this definition measures the distance from the points x and not from the sites! If we reverse the direction of measuring the distance, then we obtain a different Voronoi diagram: that obtained with the former definition using a reflected copy of P .

Many works on Voronoi diagrams based on a metric make use of the following three properties of metrics:

1. The induced topology is the same as that induced by the Euclidean metric. That is, each small neighborhood (with respect to the metric) of a point contains an L_2 -neighborhood of it, and vice versa.
2. The distance between every pair of points is invariant under translations.
3. Distances are additive along every straight line.

The convex polygon-offset distance function has properties 1 and 2 above too. However, as shown in the previous section, it does not fulfill property 3. We thus replace additivity by a weaker monotonicity property:

- 3'. Distances are monotonically increasing along every straight line.

We later use property 1 explicitly and properties 2 and 3' implicitly in the proofs of further claims.

Theorem 3.

- (a) \mathcal{D}_P has the properties 1, 2, and 3' above.
- (b) The offset bisecting curve of every pair of points in the plane is a polyline with at most $\Theta(m)$ line segments.
- (c) Two different bisecting curves intersect in at most $\Theta(m)$ points.

Proof. (a) Properties 1 and 3' follow from the fact that $O_{P,\varepsilon_1} \subseteq \text{Int}(O_{P,\varepsilon_2})$ (where $\text{Int}(P)$ denotes the interior of a polygon P) for every $\varepsilon_0 \leq \varepsilon_1 \leq \varepsilon_2$. Property 2 follows immediately from the definition of offset polygons.

(b) Assume first that the line joining the two points p, q is not parallel to any edge of P . In this case the bisecting curve of the points is a simple polyline. To observe this, pump up (offset) copies of P centered at p and q . Once the two copies make contact, two contact points continue to move away from each other. (There are always at most two such intersection points between two translated copies of the same convex

polygon.) While moving in the plane, the contact points also slide monotonically around the continually growing offset polygon. Occasionally there is a change of one of the two polygon edges that make the contact. Segments of the bisector are formed between any two such successive events. Once a polygon edge ceases to make the contact, it will never be a contact edge again. To see this consider only the edges of that copy of P centered at p . (A symmetric argument holds for the copy centered at q .) Now consider two rays emanating from p and passing through the two contact points between the two copies of P . As the copies of P grow (offset) outward, the two rays rotate monotonically in opposite directions around p , increasing the distance between the contact points and the midpoint of the segment $[pq]$. The distance between the contact points and p increases as well. Since cells of the plane subdivision induced by the medial axis of P are always opening away from the center of P (see Figure 2(a)), this trajectory cannot make the contact point enter the same cell twice. Therefore the order of formation of bisector segments resembles that of some merging of two ordered lists, each of cardinality $O(m)$. Hence the complexity of the bisector is also $O(m)$, and it is $\Theta(m)$ in the worst case.

In case there is an edge of P parallel to the line joining the two points, the bisecting curve bifurcates and the entire region bounded by the two branches of the curve (which can never intersect) is equidistant from the two points. This happens because for some values of ε there is a *segment* in the plane which is entirely of some distance δ from both points. It is easy to verify that the overall complexity of the boundary of the bisecting curve is still $O(m)$.

(c) The bisecting curve between two points is a polygonal chain, which is monotone with respect to ℓ_{pq} , the direction perpendicular to the line segment $[pq]$. Consider two pairs of points (p, q) and (r, s) . Let ξ be the direction of the bisector of the smaller angle formed by ℓ_{pq} and ℓ_{rs} . It is easy to verify that all the segments of the bisecting curve of (p, q) (resp., (r, s)) form with ℓ_{pq} (resp., ℓ_{rs}) angles of at most $\pi/4$. Therefore, the bisecting curves of (p, q) and (r, s) are both monotone with respect to ξ . Since the complexity of the two bisecting curves is $\Theta(m)$ in the worst case, the two curves intersect in at most $\Theta(m)$ points. \square

The proof of Theorem 3(b) handles the degenerate case in which \mathcal{D}_P is measured between two points which define a line parallel to some edge of P . In this special case, the bisecting curve is not a simple polyline, as explained above. The results of [KW] can still be applied with this weak form of a bisecting curve, but for simplicity of exposition we assume in the rest of the paper that the points are in general position in this respect. That is, we assume that no two input points define a line parallel to any edge of the underlying polygon.

We next follow the approach of [KW] and show some additional properties of a polygon-offset distance function and the respective nearest-site Voronoi diagram. By $|\mathcal{V}_P^N|$ we mean the total number of vertices, edges, and faces in the diagram.

Theorem 4 [KW, p. 284, Theorem 2.5]. *Assume that a distance function induces the Euclidean topology in the plane. Furthermore, assume that each bisector between two points consists of disjoint simple curves, and curves belonging to different bisectors can intersect only finitely often within each bounded area. Finally, assume that all possi-*

ble Voronoi regions are connected sets. Then $\mathcal{V}_P^N(S)$ has n faces and $O(n)$ edges and vertices.

Corollary 5. Every cell of \mathcal{V}_P^N is connected and $|\mathcal{V}_P^N| = O(nm)$.

Proof. The topology induced by \mathcal{D}_P is the same as that of the Euclidean metric (by Theorem 3(a)). Bisectors of pairs of points are composed of simple curves and intersect finitely often (by Theorem 3(b,c)). This is sufficient for applying Theorem 4 and obtaining the connectivity claim and $O(n)$ cell-complexity for a fixed value of m . The m factor is due to Theorem 3(c). □

A distance function is called *complete* if it is uniquely defined for every ordered pair of points. The next theorem establishes an important intermediate-position property for the convex polygon-offset distance function.

Theorem 6. \mathcal{D}_P is complete and for every pair of points p, q in the plane there exists a point $r \notin \{p, q\}$ such that $\mathcal{D}_P(p, r) + \mathcal{D}_P(r, q) = \mathcal{D}_P(p, q)$.

Proof. Center the polygon P at p . Sweeping the entire plane with $O_{P,\varepsilon}$ by ranging ε from ε_0 to $+\infty$ we can find the unique ε_1 for which $q \in \partial O_{P,\varepsilon_1}$. Hence \mathcal{D}_P is complete. For the second property we observe that if the entire segment $[pq]$ is dominated by some edge $e_i \in P$ (where the normal to e_i and \vec{pq} form the angle θ_i), then all the points $r \in [pq]$ fulfill the claimed equality (since then \mathcal{D}_P is merely the Euclidean distance multiplied by a constant— $\cos(\theta_i)$). This is not the case when the segment $[pq]$ is dominated by more than one edge of P . Assume that $[pq]$ is first (at the vicinity of p) dominated by the edge e_i and finally (at the vicinity of q) dominated by e_j . We have already shown in Section 2.3 that $\cos(\theta_i) < \cos(\theta_j)$ and that there exists a point $s \in [pq]$ near q such that $\mathcal{D}_P(p, s) + \mathcal{D}_P(s, q) < \mathcal{D}_P(p, q)$. It is fairly easy to find another point s' in the plane for which $\mathcal{D}_P(p, s') + \mathcal{D}_P(s', q) > \mathcal{D}_P(p, q)$: choose, for example, a point s' whose distance from \overline{pq} is twice the diameter of P . Now define a function $f(x) = \mathcal{D}_P(p, x) + \mathcal{D}_P(x, q)$. Since \mathcal{D}_P is a continuous distance function, the function $f(x)$ is continuous too, in particular along every path. Choose a path from s to s' which does not pass through p or q , e.g., the line segment $[ss']$. According to the mean-value theorem there exists a point r along this path for which $f(r) = \mathcal{D}_P(p, q)$, and the claim follows.³ □

Note that the theorem does not require r to belong to the open segment (pq) . The next theorem establishes a crucial extension property.

³ We can also find constructively such a point r . Let $d = E(p, q)/2$. Define P' (resp., P'') to be the offset of P that contains all the points r' (resp., r'') for which $\mathcal{D}_P(p, r') = d$ (resp., $\mathcal{D}_P(r'', q) = d$). By Theorem 1, the polygons P' and P'' either touch (at the midpoint of the segment $[pq]$) or intersect, thus defining the desired point r .

Theorem 7. *For every pair of points p, q in the plane there exists a point $r \neq q$ such that $\mathcal{D}_P(p, q) + \mathcal{D}_P(q, r) = \mathcal{D}_P(p, r)$.*

Proof. We use an argument similar to that of the proof of Theorem 6. As before, if the entire segment $[pq]$ is dominated by some edge $e_i \in P$ (with the respective angle θ_i between the normal to e_i and \overrightarrow{pq}), then all the points r on the extension of $[pq]$ (beyond q), which are still dominated by e_i , fulfill the claimed equality. (The case where e_i ceases to dominate \overrightarrow{pq} precisely at q is also straightforward.) Again, this is not the case when $[pq]$ is dominated by more than one edge of P . Let $[pq]$ be dominated by e_i at the vicinity of p and by e_j at the vicinity of q . By an argument similar to that of the proof of Theorem 1 we can extend infinitesimally $[pq]$ (beyond q) and find a point s on the extension for which $\mathcal{D}_P(p, q) + \mathcal{D}_P(q, s) < \mathcal{D}_P(p, s)$. This is true even if another change of the dominating edge occurs in q since \mathcal{D}_P is continuous. Hence we have $\mathcal{D}_P(p, s) - \mathcal{D}_P(q, s) > \mathcal{D}_P(p, q)$. Similarly to the previous proof, we pick another point s' in the plane for which $\mathcal{D}_P(p, q) + \mathcal{D}_P(q, s') > \mathcal{D}_P(p, s')$ (this time choose a point s' very close to p but not on \overrightarrow{pq}). Hence we have $\mathcal{D}_P(p, s') - \mathcal{D}_P(q, s') < \mathcal{D}_P(p, q)$. Now define a function $g(x) = \mathcal{D}_P(p, x) - \mathcal{D}_P(q, x)$. Since \mathcal{D}_P is a continuous distance function, the function $g(x)$ is continuous too, in particular along every path. Choose a path from s to s' which does not pass through q , e.g., the line segment $[ss']$. According to the mean-value theorem there exists a point r along this path for which $g(r) = \mathcal{D}_P(p, q)$. Evidently $r \neq q$, and the claim follows.⁴ \square

Theorem 8 [KW, p. 286, Theorem 4.1]. *Assume that a distance function induces the Euclidean topology in the plane. Furthermore, assume that the distance function fulfills Theorems 6 and 7. Then all Voronoi regions are simply connected.*

Corollary 9. *Every cell of \mathcal{V}_P^N is simply connected.*

Proof. The topology induced by \mathcal{D}_P is the same as that of the Euclidean metric (by Theorem 3(a)). \mathcal{D}_P is complete, and for every pair of points p, q in the plane there exists a point $r_1 \notin \{p, q\}$ such that $\mathcal{D}_P(p, r_1) + \mathcal{D}_P(r_1, q) = \mathcal{D}_P(p, q)$ (by Theorem 6) and a point $r_2 \neq q$ such that $\mathcal{D}_P(p, q) + \mathcal{D}_P(q, r_2) = \mathcal{D}_P(p, r_2)$ (by Theorem 7). This is sufficient for applying Theorem 8 and obtaining the claim. \square

3.2. Computing the Diagram

For computing Voronoi diagrams based on the polygon-offset distance function we need two primitive functions: one that computes the offset distance between two given points, and another that computes the Voronoi vertex of three given points.

⁴ Similarly to the previous theorem, we can also find constructively such a point r by defining $d = E(p, q)$ and by intersecting the polygons $\{r' \mid \mathcal{D}_P(p, r') = 2d\}$ and $\{r'' \mid \mathcal{D}_P(q, r'') = d\}$.

Theorem 10. *Allowing $O(m)$ -time preprocessing of P , $\mathcal{D}_P(p, q)$ can be computed in $O(\log m)$ time for every pair of points p, q in the plane.*

Proof. The procedure for using a polygon P for measuring a convex (scaled) distance function between points in the plane is rather simple. A *binary search* is performed (in $O(\log m)$ time) in the cyclic ordered list of edge-slopes of P for finding the edge e that dominates \vec{pq} at q . Then the distance is easily determined.

For a polygon-offset distance function we preprocess P in $O(m)$ time and build an appropriate (tree-like) data structure that represents the medial axis of P . Each segment of the data structure represents a vertex (or several consecutive vertices) of P and each region is attributed by the dominating edge. The data structure answers, given a vector \vec{pq} positioned such that p coincides with the center of P , in which region of the medial axis q falls. This region defines the edge of P that dominates q . Locating this region requires $O(\log m)$ time. The offset distance from p (now located at the center of P) to the medial-axis vertex (that is the root of the located region) requires constant time. (This information is stored at the medial-axis vertices during the preprocessing.) Finally, the offset distance from the medial-axis vertex to q is computed in constant time using simple algebra. Thus, the entire computation requires $O(\log m)$ time. \square

For the second primitive function we need the prune-and-search technique of Kirkpatrick and Snoeyink [KS]. We provide here a brief description of this technique. The goal is to find efficiently a fixed point of the composition of two or three monotone and continuous functions that are defined piecewise. Such a function is defined in some range, which is partitioned into a set of closed pairwise-disjoint (except in their endpoints) intervals, whose union is the original range. In each of these intervals the function can be evaluated in constant time. Such a function is therefore called *piecewise-basic*. It is proven in [KS] that given three monotone-*decreasing* piecewise-basic functions, which are defined on the same range and with m intervals each, the fixed point of their composition can be computed in optimal $\Theta(\log m)$ time. This is a significant improvement over the commonly used $O(\log^2 m)$ -time nested binary search. The tentative prune-and-search technique maintains for each function three types of intervals: those in which the fixed point cannot lie (eliminated), those in which the sought point probably does not lie (tentatively pruned out), and those in which the point may lie (in doubt). During the course of the algorithm intervals are tentatively discarded; these discards are later revoked or certified. The running time of the procedure follows from the fact that at each step a fixed fraction of the doubted intervals are eliminated or tentatively pruned out (with the assurance that a fixed fraction of the discards are correct), or a fixed fraction of the tentative discards are revoked or certified. The procedure terminates when each of the three functions is restricted to only one interval. At this point a constant-time step is invoked to find the actual fixed point.

The technique is applied in proving the following theorem:

Theorem 11 [KS, Section 3.4, Theorem 3.10]. *Given three convex polygons P, Q, R of total complexity m , one can determine the point that is equidistant from P, Q, R in $O(\log m)$ time by tentative prune-and-search.*

The proof of the above theorem starts with restricting the attention to three convex chains, subsets of P , Q , R , for which the Voronoi vertex problem has a unique solution. Then the polygons are discretized, parameterized clockwise in the range $[0, 1]$, and intervals are defined according to the polygon vertices. Appropriate monotone-decreasing functions are then defined on the polygons, such that their Voronoi vertex is nothing else but the fixed point of the composition of the three functions. More specifically, the function $f_P(p')$ is defined to be the parameter of the point $q' \in Q$ whose normal intersects the normal to P at p' in a point v such that $|p'v| = |q'v|$. Similarly, $f_Q(q')$ and $f_R(r')$ are defined by using R and P , respectively. At each step of the algorithm a statement of the form $f_P(p') > q'$ is asserted, guiding discards of intervals (permanently or tentatively), or the certification or revoking of previous discards. When all three functions are restricted to a single interval (polygon edge), simple algebra is used to determine the actual Voronoi vertex.

Theorem 12. *Given three points p, q, r in the plane and a convex m -edge polygon P , we can find the point equidistant from p, q, r (according to \mathcal{D}_P) in $O(\log^2 m)$ time (allowing $O(m)$ time for preprocessing P).*

Proof. First we need to argue that the point sought exists and is unique. This point is the intersection of the bisector of p and q and the bisector of p and r . As was shown earlier in the paper, these bisectors are monotone with respect to the perpendicular bisectors of the segments $[p, q]$ and $[p, r]$, respectively. Thus they intersect exactly once at the desired Voronoi vertex.

We apply a prune-and-search procedure similar to that used in the proof of Theorem 11 for locating the point equidistant (with respect to the regular Euclidean distance function) from three given convex polygons.

The main difficulty that we face is carrying the polygon discretization of the problem in [KS] over to the computation of the point offset-wise equidistant from three given points. While the curve discretization of [KS] is static (as required for applying prune-and-search), it seems that in our problem the discretization should be dynamic. This is because the front of the offset polygon undergoes topological changes: new vertices appear when the polygon is offset outward from its center. Nevertheless, we can still use static discretization. Instead of discretizing a curve, we discretize the circle of directions centered at the center of the polygon P . The discrete set of directions consists of all the rays emanating from p pointing at vertices of P and to vertices of its medial axis ($2m - 3$ directions in total); see Fig. 7. We need not worry that at some (small) offset distances some of the polygon edges do not “exist” yet. (That is, that halfplanes that correspond to offset edges are redundant.) This is taken care of by the primitive function that computes offset distances between points.

The crucial observation is that we can apply the same decision paradigm as in Theorem 3.10 of [KS]. As in [KS], we shoot straight rays from the points p and q in some candidate directions p' and q' , and compute their intersection point v . However, we use the offset distance function for comparisons, that is, we compare $\mathcal{D}_P(p', v)$ to $\mathcal{D}_P(q', v)$. Since the offset distance is monotonically increasing with respect to the Euclidean distance, we may apply the same paradigm, discarding (either permanently or tentatively) half of the candidate directions for some of the three input points.




Fig. 7. Discretization of the offset function for tentative prune-and-search.

The search procedure terminates when we identify the intervals e_p, e_q, e_r (edges of the underlying polygon P centered at p, q, r , respectively) that contain the exact directions p^*, q^*, r^* from the points p, q, r , respectively, that define the Voronoi vertex. Unfortunately, within these intervals the distance function is not basic (see Fig. 7). Hence, unlike the problem of [KS], p^*, q^*, r^* cannot be determined in additional constant time by simple algebra. In the worst case each of e_p, e_q, e_r is dominated (at different offset distances) by as many as $\Theta(m)$ polygon edges. The point offset-wise equidistant from the three input points is found in the intersection of three slabs, $\sigma_p, \sigma_q, \sigma_r$, emanating from p, q, r , respectively.

We now apply a secondary, this time regular, prune-and-search step in order to locate the desired point. The discretization of this subproblem is along the slabs (away from p, q, r) according to their contained regions. Given three candidate regions ρ_p, ρ_q, ρ_r in the slabs $\sigma_p, \sigma_q, \sigma_r$, respectively, we choose from each region the offset-wise closest point to p, q, r , respectively, as the representative points, compare the offset distances, and apply the regular tentative prune-and-search technique. This step computes the exact regions in the slabs in which the desired Voronoi vertex is located. When these three regions, one in each slab, are known, simple algebra completes the calculation.

The running time of the main tentative prune-and-search step is $O(\log^2 m)$. It is obtained similarly to that of the problem of [KS], with the only difference that the $O(1)$ -time (Euclidean) distance-evaluation function of that problem is replaced by an $O(\log m)$ -time function for computing the offset distance (Theorem 10). The secondary prune-and-search step requires $O(\log m)$ time, and the completion step requires only constant time. Overall, the entire procedure requires $O(\log^2 m)$ time. \square

We now define (following [KW]) the $(\alpha, \bar{\alpha})$ -support as follows. A distance function D has an $(\alpha, \bar{\alpha})$ -support (for $\alpha, \bar{\alpha} \in [0, \pi], \alpha \neq \bar{\alpha}$) if it fulfills the following condition for any pair of points p, q in the plane, where \overline{pq} is a line of slope α : if ℓ is the line of slope $\bar{\alpha}$ that passes through q , then all points r that satisfy $D(p, r) \leq D(p, q)$ lie on the same side of ℓ as p (see Fig. 8).

Lemma 13. *There exist angles $\alpha \neq \beta$ and $\bar{\alpha}, \bar{\beta}$ such that \mathcal{D}_P has both $(\alpha, \bar{\alpha})$ - and $(\beta, \bar{\beta})$ -supports.*




Fig. 8. $(\alpha, \bar{\alpha})$ -support.

Proof. There exist three edges $e_i, e_j, e_k \in P$ that never disappear in inner offsets of P (in the degenerate case there are only two such parallel edges). Without loss of generality assume that the normal vectors to e_i and e_j have slopes in $[0, \pi]$. Consider the edge e_i and the normal to it ℓ_i . Set α to the slope of ℓ_i and $\bar{\alpha}$ to the slope of e_i . It is trivial to verify that \mathcal{D}_P has an $(\alpha, \bar{\alpha})$ -support. Similarly set β to the slope of ℓ_j and $\bar{\beta}$ to the slope of e_j . \mathcal{D}_P also has a $(\beta, \bar{\beta})$ -support. Obviously $\alpha \neq \beta$ and $\bar{\alpha} \neq \bar{\beta}$, establishing the claim. \square

This allows us to apply the algorithmic framework of [KW]:

Theorem 14 [KW, p. 288, Theorem 5.1]. *Assume that a distance function induces the Euclidean topology in the plane. Furthermore, assume that each bisector between two points consists of disjoint simple curves, and curves belonging to different bisectors can intersect only finitely often within each bounded area. Also assume that the distance function fulfills Theorem 6. Finally, assume that there are angles α, β , and $\bar{\alpha} \neq \bar{\beta}$, such that the distance function has both $(\alpha, \bar{\alpha})$ - and $(\beta, \bar{\beta})$ -support. Then $\mathcal{V}_P^N(S)$ can be computed within optimal $O(n \log n)$ steps.*

Corollary 15. *Let S be a set of n points in the plane. A compact representation of $\mathcal{V}_P^N(S)$ can be computed in expected $O(n \log n \log^2 m + m)$ time.*

Proof. The topology induced by \mathcal{D}_P is the same as that of the Euclidean metric (by Theorem 3(a)). Bisecting curves of pairs of points are simple and intersect finitely often (by Theorem 3(b,c)). \mathcal{D}_P is complete and for every pair of points p, q in the plane there exists a point $r \notin \{p, q\}$ such that $\mathcal{D}_P(p, r) + \mathcal{D}_P(r, q) = \mathcal{D}_P(p, q)$ (by Theorem 6). There exist angles $\alpha \neq \beta$ such that \mathcal{D}_P has both $(\alpha, \bar{\alpha})$ - and $(\beta, \bar{\beta})$ -supports (by Lemma 13). This is sufficient for applying Theorem 14 and obtaining the claim for a constant value of m . The extra $\log^2 m$ factor is due to the primitive function that computes a Voronoi vertex of three given points (see Theorem 12). The algorithm calls this function $O(n \log n)$ times (due to the randomized divide-and-conquer paradigm). Bisecting curves (Voronoi edges connecting Voronoi vertices) are then represented implicitly by the relative positions of the sites that define them. \square

We can achieve a faster running time than this, however.

4. Compact Nearest-Site Voronoi Diagram

We now extend the work of McAllister et al. [MKS] to polygon-offset distance functions. They describe in detail an algorithm for constructing a *compact* nearest-site Voronoi diagram, where the underlying convex distance function is defined by a convex polygon. (They also consider sites as being convex polygons; we need only consider point sites.)

The compact diagram simplifies the full Voronoi diagram by maintaining a coarse “dual” of it: for each vertex of the full diagram, the compact diagram maintains a set of *spokes* (minimum-length segments from the vertex to the polygonal sites, according to the convex distance function), and each polygonal site is replaced by the convex hull of vertices of the site in which spokes occur. (See Fig. 4 on p. 81 of [MKS] for an illustration.) This allows the complexity of the compact diagram to be $O(n)$ instead of $O(nm)$, where m is the complexity of P (the underlying convex polygon that defines the convex distance function), and n is the number of points.

The compact diagram does not allow definite point location queries, as each cell of it is composed of portions of two cells of the full Voronoi diagram. However, with a low cost of $O(\log n + \log m)$ time the compact diagram provides two candidates for the point location problem. As the authors note, this is sufficient for the *post-office* problem and for *retraction motion planning*. The compact diagram reduces two components of the full Voronoi diagram: the descriptions of the sites and the descriptions of the bisecting curves between the sites. (Since the sites are not points and the distance function is not Euclidean, a bisecting curve of two sites is a polyline instead of a line.) We deal with point sites and so we do not benefit from the first reduction; still our distance function is not Euclidean and hence we benefit from the second reduction.

We assume the familiarity of the reader with [MKS] and proceed to sketch how to generalize the compact Voronoi diagram to a convex polygon-offset distance function. In order to do that we have to make sure that the geometric properties of the compact Voronoi diagram are preserved when we change the distance function from convex to convex polygon-offset. For this goal we need only observe that for every $\varepsilon_1 > \varepsilon_2$ the convex polygon O_{P,ε_1} fully contains the convex polygon O_{P,ε_2} and that two offsets of the same polygon intersect at most twice [BBDG, Theorem 1]. (The second property is crucial for applying the *tentative prune-and-search* technique of [KS].) This suffices to prove the spoke properties (Lemmas 2.2–2.6 of [MKS]) and the correctness of the plane-sweep algorithm [ibid., Lemmas 3.1–3.4 and Theorem 3.5].

Thus we are able to replicate the claims of [MKS] regarding the running time of the primitive operations (Lemmas 3.13–3.15). In particular, the functions *spoke* (p, A) and *vertex* (ABC) require $O(\log^2 m)$ time. The first function computes the spoke that connects the point p (where the center of the m -gon is positioned) to a site A , and the second function computes the Voronoi vertex (finite or infinite) around which three given sites A, B, C appear in counterclockwise order. These functions are implemented in [MKS] so that they require $O(\log n + \log m)$ and $O(\log n \log m)$ time, respectively, but they run in only $T_v = O(\log m)$ time in the special case where all the sites are points instead of convex polygons. The vanishing of the $\log n$ factor in this case carries over also to convex polygon-offset distance functions. However, the $\log m$ term becomes $\log^2 m$ by applying the procedure of Theorem 12. In conjunction, Theorem 3.10 on p. 94 of

[MKS] states that the compact representation of the Voronoi diagram can be computed in worst-case $O(n(\log n + T_v))$ time, resulting in a near-optimal $O(n(\log n + \log^2 m) + m)$ -time algorithm for computing a compact representation of \mathcal{V}_P^N . (Obviously, $\Omega(n(\log n + \log m) + m)$ is a lower bound on the running time of any such algorithm. Had we been able to compute the Voronoi vertex of three points in $O(\log m)$ time, we would have obtained an optimal-time algorithm.)

5. Furthest-Site Voronoi Diagram

In this section we show how to construct the furthest-site Voronoi diagram with respect to \mathcal{D}_P . For this purpose we show that \mathcal{D}_P fits the framework of Mehlhorn et al. [MMR] which follows Klein's unifying approach for Voronoi diagrams [KW], [KJ]. Mehlhorn et al. show that under some conditions the complexity of the furthest-site Voronoi diagram is $O(n)$, where n is the number of sites.

First we adopt some terminology of [MMR]. Let p, q be a pair of points in the plane. The *dominant set* $M_P(p, q)$ (hereafter denoted as $M(p, q)$) contains all the points that are closer to p than to q with respect to \mathcal{D}_P . Let $S = \{p_i | 1 \leq i \leq n\} \in \mathbb{R}^2$ be a set of n points. The family $\mathcal{M} = \{M(p_i, p_j) | 1 \leq i \neq j \leq n\}$ is called a *dominance system* if for all $p_i, p_j \in S$ we have:

1. $M(p_i, p_j)$ is open and nonempty;
2. $M(p_i, p_j) \cap M(p_j, p_i) = \emptyset$ and $\partial M(p_i, p_j) = \partial M(p_j, p_i)$; and
3. $\partial M(p_i, p_j)$ is homeomorphic to the open interval $(0, 1)$.

Theorem 16. \mathcal{M}_P (the family \mathcal{M} with respect to \mathcal{D}_P) is a dominance system.

Proof. The dominance-system properties of \mathcal{M}_P follow immediately from the definition of \mathcal{D}_P :

1. $M(p_i, p_j)$ is open because \mathcal{D}_P is monotone and continuous. It is obviously nonempty: assume without loss of generality that the line that connects p_i and p_j is horizontal and that p_i is on the left of p_j . Then all the points on the left of p_i (resp., right of p_j) belong to $M(p_i, p_j)$ (resp., $M(p_j, p_i)$) regardless of P .
2. Although \mathcal{D}_P is not symmetric, the bisecting curves of the pairs of points (p_i, p_j) and (p_j, p_i) are obviously identical. Centering P at any point r and "pumping" P up, it either hits p_i first, p_j first, or simultaneously p_i and p_j . (This is because \mathcal{D}_P is continuous and monotonically increasing.) In the first case $r \in M(p_i, p_j)$, in the second $r \in M(p_j, p_i)$, and in the third $r \in \partial M(p_i, p_j) = \partial M(p_j, p_i)$.
3. This follows from the fact that $\partial M(p_i, p_j)$ is a simple polygonal chain. □

Assume that for the nearest-site Voronoi diagram every portion of the bisecting curve $\partial M(p_i, p_j)$ is put in the cell of $\min(i, j)$.

A dominance system is called *admissible* if it satisfies in addition the conditions:

4. Bisectors intersect finitely many times.

5. For all $S' \subset S$, $S' \neq \emptyset$ and for every reordering of indices of points in S :
 - (a) The Voronoi cell of every point $p_i \in S'$ (with respect to S') is connected and has a nonempty interior.
 - (b) The union of all the Voronoi cells of all points $p_i \in S'$ (with respect to S') is the entire plane.

A dominance system which fulfills only properties 4 and 5b is called *semi-admissible*.

Theorem 17. \mathcal{M}_P is admissible.

The proof of Theorem 17 is straightforward (part of the theorem is proven earlier in this paper) and is omitted here.

We now consider \mathcal{M}^* , the “dual” of \mathcal{M} , in which the dominance relation as well as the ordering of the points are reversed.

Theorem 18 [MMR, Lemma 1]. *If \mathcal{M} is semi-admissible then so is \mathcal{M}^* . Moreover, the (so-called furthest-site) Voronoi diagram that corresponds to \mathcal{M}^* is identical to the (nearest-site) Voronoi diagram that corresponds to \mathcal{M} .*

Note that admissibility is not preserved when moving to the dual of the dominance system. This corresponds to the fact that cells in the furthest-site Voronoi diagram may well be disconnected.

We have thus shown that the furthest-site Voronoi diagram (with respect to \mathcal{D}_P) can be defined in terms of a dominance system (by moving to its dual). Therefore the results of [KI], [KW], and [KMM] apply to both nearest- and furthest-site Voronoi diagrams. Furthermore, we can benefit from all the results of [MMR] on this diagram, namely, that it is a tree, and that we can compute it by a randomized algorithm in $O(n \log n \cdot K + m)$ time, where K is the time needed by a primitive function that considers the diagram of five sites. Applying the same technique of Theorem 12, this primitive requires $O(\log^2 m)$ time in our setting. Hence the overall expected time required for computing \mathcal{V}_P^F is $O(n \log n \log^2 m + m)$. As with the nearest-site diagram, we can do better than this.

We have followed the above approach of [KMM] only for proving the properties of the furthest-site Voronoi diagram. Our real goal is to obtain a deterministic algorithm for computing the diagram. For this purpose we adopt the three-dimensional plane-sweep approach of Rappaport [Ra] which follows Fortune’s algorithm [Fo]. In this approach axis-parallel cones are emanating from the sites and the plane-sweep detects their intersections and produces the corresponding Voronoi vertices. The major detail that we need to note is that in our setting we have polyhedral cones, and that we can compute the intersection of any three such cones in $O(\log^2 m)$ time. (The xy -projections of these intersection points are the Voronoi vertices.) Thus we spend $O(n \log n + m)$ time in preprocessing and $O(n(\log n + \log^2 m))$ time in the sweep (n events, for each we spend $O(\log^2 m)$ time for computing and $O(\log n)$ time for queue operations), resulting in a near-optimal $O(n(\log n + \log^2 m) + m)$ -time algorithm. We again store \mathcal{V}_P^F compactly, representing bisecting curves implicitly via the relative positions of the two sites that define them.

6. Conclusion

We develop in this paper the notion of a convex polygon-offset distance function. This is an extremely important notion for tolerancing in manufacturing. We describe in detail how to compute the nearest- and furthest-site Voronoi diagrams with respect to this distance function. In a companion paper [BBDG] we use these diagrams for solving a tolerancing problem: given a set S of points in the plane and a convex polygon P , find the minimum ε and the respective translation τ for which the ε -offset annulus of $\tau(P)$ covers S .

For simplicity of exposition we have defined the offset so as to be piecewise-linear. However, this can be extended to the usual definition of a true offset, where every point on the outer offset is within the same distance from the original polygon. In this case the outer offset is made up of alternating line segments and circular arcs. The latter offset definition captures better tolerancing problems that arise in manufacturing, robot motion planning, etc. [BBDG]. It is easy to verify that all the claims in this paper still hold in the true-offset case, and that the algorithms require asymptotically the same amount of time. This is mainly because (1) the complexity of an outer offset polygon is only twice the complexity of the original polygon; (2) like segment–segment intersections, computing arc–arc and arc–segment intersections can also be performed in constant time; and (3) the number of bisector–bisector intersections is asymptotically the same with respect to the two offset definitions.

Acknowledgments

We wish to thank J. Snoeyink and P. Bose for helpful discussions on Voronoi diagrams and methods for computing them.

References

- [AA] O. Aichholzer and F. Aurenhammer, Straight skeletons for general polygonal figures in the plane, *Proc. 2nd COCOON*, Lecture Notes in Computer Science, 1090, Springer-Verlag, Berlin, 1996, pp. 117–126.
- [AAAG] O. Aichholzer, D. Albers, F. Aurenhammer, and B. Gärtner, A novel type of skeleton for polygons, *Journal of Universal Computer Science* (an electronic journal), **1** (1995), 752–761.
- [BBDG] G. Barequet, A. Briggs, M. Dickerson, and M. T. Goodrich, Offset-polygon annulus placement problems, *Computational Geometry: Theory and Applications*, **11** (1998), 125–141.
- [CD] L. P. Chew and R. L. Drysdale, Voronoi Diagrams Based on Convex Distance Functions, Technical Report PCS-TR86-132, Dept. of Computer Science, Dartmouth College, Hanover, NH 03755, 1986; a preliminary version appeared in *Proc. 1st Ann. ACM Symp. on Computational Geometry*, Baltimore, MD, 1985, pp. 235–244.
- [Cr] C. M. Creveling, *Tolerance Design: A Handbook for Developing Optimal Specifications*, Addison-Wesley, Reading, MA, 1997.
- [Du] J. Dugundji, *Topology*, Allyn and Bacon, Boston, MA, 1970.
- [Fo] S. Fortune, A sweepline algorithm for Voronoi diagrams, *Algorithmica*, **2** (1987), 153–174.
- [Kl] R. Klein, *Concrete and Abstract Voronoi Diagrams*, Lecture Notes in Computer Science, 400, Springer-Verlag, Berlin, 1989.
- [KMM] R. Klein, K. Mehlhorn, and S. Meiser, Randomized incremental construction of abstract Voronoi diagrams, *Computational Geometry: Theory and Applications*, **3** (1993), 157–184.

- [KN] J. L. Kelley and I. Namioka, *Linear Topological Spaces*, Springer-Verlag, New York, 1976.
- [KS] D. Kirkpatrick and J. Snoeyink, Tentative prune-and-search for computing fixed-points with applications to geometric computation, *Fundamental Informaticæ*, **22** (1995), 353–370.
- [KW] R. Klein and D. Wood, Voronoi diagrams based on general metrics in the plane, *Proc. 5th Symp. on Theoretical Aspects in Computer Science*, Lecture Notes in Computer Science, 294, Springer-Verlag, Berlin, 1998, pp. 281–291.
- [MKS] M. McAllister, D. Kirkpatrick, and J. Snoeyink, A compact piecewise-linear Voronoi diagram for convex sites in the plane, *Discrete & Computational Geometry*, **15** (1996), 73–105.
- [MMO] K. Mehlhorn, S. Meiser, and Ó'Dúnlaing, On the construction of abstract Voronoi diagrams, *Discrete & Computational Geometry*, **6** (1991), 211–224.
- [MMR] K. Mehlhorn, S. Meiser, and R. Rasch, Furthest Site Abstract Voronoi Diagrams, Technical Report MPI-I-92-135, Max-Planck-Institut für Informatik, Saarbrücken, 1992.
- [OBS] A. Okabe, B. Boots, and K. Sugihara, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Wiley, Chichester, 1992.
- [Or] J. O'Rourke, *Computational Geometry in C*, Cambridge University Press, New York, 1994.
- [Ra] D. RAPPAPORT, Computing the furthest site Voronoi diagram for a set of disks, *Proc. 1st Workshop on Algorithms and Data Structures*, Lecture Notes in Computer Science, 382, Springer-Verlag, Berlin, pp. 57–66, 1989.

Received January 8, 1999, and in revised form January 5, 2000, and June 12, 2000.

Online publication December 4, 2000.