Voronoi Diagrams for Polygon-Offset Distance Functions*

Gill Barequet[1], Matthew T. Dickerson[2], and Michael T. Goodrich[1]

[1] Center for Geometric Computing, Dept. of Computer Science, Johns Hopkins University, Baltimore, MD 21218. E-mail: `[barequet|goodrich]@cs.jhu.edu`

[2] Department of Mathematics and Computer Science, Middlebury College, Middlebury, VT 05753. E-mail: `dickerso@middlebury.edu`

Abstract. In this paper we develop the concept of a *polygon-offset* distance function and show how to compute the respective nearest- and furthest-site Voronoi diagrams of point sites in the plane. We provide optimal deterministic $O(n(\log n + \log m) + m)$-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.

Keywords: Voronoi diagrams, medial axis, distance function, offset, convexity, geometric tolerancing.

1 Introduction

The Voronoi diagram of a set $S \subset \mathbb{R}^2$ is a powerful tool for handling many geometric problems dealing with distance relationships. Voronoi diagrams have been used extensively, for example, for solving nearest-neighbor, furthest-neighbor, and matching problems in many contexts. The underlying distance function 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 distance functions, which are extensions of the scaling notion for circles to convex polygons. Nevertheless, we feel that defining distance in terms of an offset from a polygon is more natural than scaling in many applications, including those dealing with manufacturing processes. This is because the relative error of the production tool is independent of the location of the produced feature relative to some artificial reference point (the "origin"). Therefore it is more likely to allow (and expect) local errors bounded by some tolerance, rather than errors scaled around some (arbitrary) center.

In this paper we investigate distance functions based on *offsetting* convex polygons, where the distance is measured along the infinitely extended medial-axis of such a polygon. 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.

---

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

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 was the first to study the related notion of *convex distance functions*. He showed, for example, that distance can be defined in terms of a scaling of a convex polygon, and that while such functions do not in general define metrics, they exactly characterize the "distance" functions satisfying the triangle inequality. Chew and Drysdale [CD] show that one can 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, as they show the Voronoi diagram can be of size $\Theta(nm)$. This work was generalized even further [Kl, KMM, KW, MMO, MMR], showing how to define Voronoi diagrams in a very abstract setting. They also give randomized incremental constructions for this abstract setting that can be applied to nearest- and furthest-neighbor Voronoi diagrams for convex distance functions. 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, Kirkpatrick, and Snoeyink [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 present in [AA] 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., Figure 4] of a Voronoi diagram of two straight segments. They take the set of polygonal chains which define the distance function to also be the set of sites for which the Voronoi diagram is sought. We, however, make the distinction between the convex polygon which defines the distance function and the set of points, 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. Nevertheless, we show that convex polygon-offset distance functions satisfy all the topological properties for abstract Voronoi diagrams [Kl, KMM, KW, MMO, MMR]. Finally, given a set S

of n points in the plane, we show how to deterministicly construct compact representations of nearest- and furthest-site Voronoi diagrams for S with respect to an offset distance defined by an m-edge convex polygon in $O(n(\log n + \log m) + m)$ time. We apply the tentative prune-and-search paradigm of Kirkpatrick and Snoeyink [KS], as utilized by McAlister et al. [MKS], to the polygon-offset distance function. Like the method of McAlister et al., our nearest-neighbor diagram is based upon using this approach to design a non-trivial generalization of Fortune's plane-sweep algorithm [Fo], whereas our furthest-neighbor diagram algorithm is based upon applying tentative prune-and-search to a generalization of Rappaport's space-sweeping algorithm [Ra].

2 Preliminaries

2.1 Convex Polygon-Offset Distance Functions

We first define the *offset* of a convex polygon. For simplicity of expression in this extended abstract we define the offset so as to be piecewise-linear, and we show in the full version that this can be extended to the usual definition of an offset (when, for example, an outer offset is made up of alternating line segments and circular arcs).

The *outer* ε-offset of P is obtained by translating each edge $e \in P$ by ε in a direction orthogonal to e and by extending it so as to meet the translations of the edges neighboring to e. The edge e' is trimmed by the lines parallel and at distance ε (outside of P) of the neighboring edges of e. The *inner* offset is defined in a similar way: For each edge $e \in P$ we construct a line parallel to it and at distance ε on the inside of P. If we increase ε continuously we observe that the edge e'' (the inner offset of e) gets smaller continuously and may even "disappear". This happens when the neighboring offset lines meet "before" they intersect with the line that corresponds to e''. Figure 1 shows inner and outer offsets of a convex polygon for some value of ε. We adopt the notation of [BBDG] and denote the inner and outer ε-offset polygons of P by $I_{P,\varepsilon}$ and $O_{P,\varepsilon}$, respectively. In addition, the notation $O_{P,-\varepsilon}$ is used as a synonym for $I_{P,\varepsilon}$, in particular when the sign of ε is unknown.

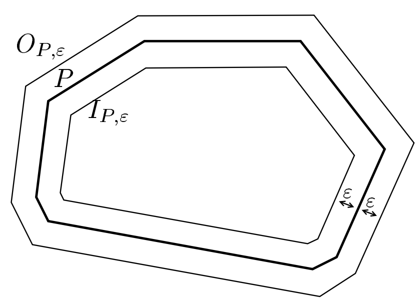

Fig. 1. Inner and outer offsets

As long as we increase ε (for inner offsets) more and more edges of $I_{P,\varepsilon}$ "disappear" in a well-defined order. By increasing the value of ε we actually move all the vertices of P along (and inward) the medial-axis of P. When two neighboring vertices of P meet at a junction of the medial axis the corresponding edge disappears and the two vertices are merged into one vertex. This process continues and at some point we are left with a triangle which in turn eventually collapses into a point for some value of ε. We denote this point as the *center* of P.

Figure 2 illustrates the medial-axis-based offset operation. In case P contains two parallel edges which also define the maximum width of P, $I_{P,\varepsilon}$ becomes (for some value of ε) a segment that can no more be offset inwards. Let ε_0 be the value of ε for which I_{P,ε_0} degenerates into a point or a segment. ε_0 is also the Euclidean distance from the center to the nearest edge of P. The outer offset behaves more "normally": no edges appear or disappear.

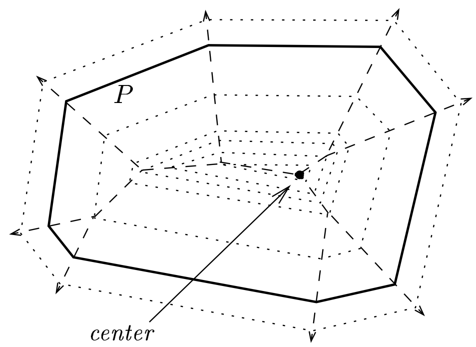

Fig. 2. Offset and the medial-axis

We now define the convex polygon-offset distance function $\mathcal{D}_P$ between two points p and q. Let d be the Euclidean distance from the center to the closest edge in the offset of P centered at p which first touches q. Then $\mathcal{D}_P(p,q) = d/\varepsilon_0$. This gives us a similar normalization where $\mathcal{D}_P(p,p) = 0$, and $\mathcal{D}_P(p,q) = 1$ if q is on the unit polygon.

The polygon-offset distance function has several interesting properties. In general, if the polygon P is not regular then the polygon-offset distance function $\mathcal{D}_P$ is not symmetric. Thus it is not a metric. Moreover, it does not even obey the triangle inequality, as we show below. Nevertheless, we show that the polygon-offset distance function $\mathcal{D}_P$ *can* be used as a basic distance function in a well-defined Voronoi diagram. We will denote the nearest- and furthest-site Voronoi diagrams with respect to $\mathcal{D}_P$ by $\mathcal{V}_P^n$ and $\mathcal{V}_P^f$, respectively.

2.2 Basic Properties

For measuring the distance from some point p to another point q we center P at p and offset it until it hits q. We say that the edge of P that hits q *dominates* it. We show later in this section how the dominating edge may change along some direction. Let us consider what happens as we move points farther from p, the center of P, in some direction defined by a ray $\overrightarrow{pq}$. When $\overrightarrow{pq}$ crosses an edge of the medial-axis, this corresponds to switching from the dominance of one edge to the dominance of another edge, as each region of the plane subdivision induced by the (infinite-extension of the) medial axis corresponds to an edge. Figure 3 illustrates this situation: The direction $\overrightarrow{pq}$ (p is positioned at the center of P) is first dominated by $e_1 \in P$ (in the inner offsets of P the edge e_2 does not exist yet), and then it crosses an edge of the medial-axis and becomes dominated by e_2.

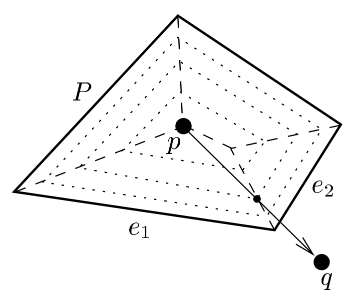

Fig. 3. Dominating edges

We now give some more specific relationships between $\mathcal{D}_P$ and the Euclidean distance E. Refer to Figure 4. Assume that the dominating edge e is moving in a direction ℓ_e from the center point p (so that ℓ_e is perpendicular to e). In moving from p to q along this "wave" e, denote the angle between ℓ_e and $\overrightarrow{pq}$ by θ. A unit move of e in the direction of $\overrightarrow{pq}$ corresponds to an offset of $\cos(\theta)$ along ℓ_e. Intuitively, if the offset from P is growing along a moving "wave" e, the point of contact is like a "surfer" moving on this wave in the direction ℓ_e. Increasing the offset by 1 increases the Euclidean distance in the direction of $\overrightarrow{pq}$ by $1/\cos\theta$. When $\theta = 0$, ℓ_e and $\overrightarrow{pq}$ have the same "speed". As the absolute value of θ increases, a unit offset of an edge outward results in an increasing Euclidean distance. Conversely, as θ increases, a unit Euclidean distance corresponds to a decreasing distance with respect to $\mathcal{D}_P$ along ℓ_e, until it vanishes at $\theta = \pi/2$. Intuitively, the "speed" along $\overrightarrow{pq}$ is inversely proportional to the charged $\mathcal{D}_P$ distance (the offset). At $\theta = \pi/2$ the charge (offset) is 0 and hence the "speed" is infinite.

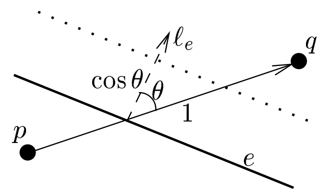

Fig. 4. $\mathcal{D}_P$ and the Euclidean distance

This explains why the "speed" along any direction v can only decrease. Consider the situation where v crosses an edge of the medial-axis. That is, it moves from a region that corresponds to some edge $e_1 \in P$ to the region that corresponds to another edge e_2. The edge of the medial-axis is the bisector of e_1 and e_2. Elementary geometry shows that v forms a larger angle with the normal to e_1 than that with the normal to e_2 and therefore its "speed" decreases, that is, the offset charge is increased. Hence we have:

Theorem 1. *If p, s, q are colinear 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 every point on the edge pq is dominated by the same edge of P when P is centered at p.*

Corollary 2. *A polygon-offset distance function does* **not** *fulfill the triangle inequality.*

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

3 Nearest-Site Voronoi Diagram

Let us now address the construction of a nearest-site Voronoi diagram. We begin by showing that this notion of distance fits the unifying approach of Klein [KW, Kl]. Klein proposes to replace the *distance* notion by *bisecting curves*. Each pair of sites are separated by curves which divide the plane into two portions each corresponding to one site. The Voronoi region of some site s is the intersection of all the s-portions defined by bisecting curves between s and all the other sites.

3.1 Properties of $\mathcal{D}_P$

In order to show that the polygon-offset distance function is valid for defining Voronoi diagrams, we need to precisely define the nearest-site Voronoi diagram in this context: The Voronoi cell that corresponds to a point p consists of all the points x in the plane for which $\mathcal{D}_P(x,p) \le \mathcal{D}_P(x,q)$ for all $q \ne p$. 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.

Klein makes use of the following three properties of *metrics*:
1. The topology induced is the same as that induced by the Euclidean metric;
2. The distance between every pair of points is invariant under translations;
3. Distances are additive along every straight line.
In fact, the third property can be replaced by a weaker property:

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

Theorem 3. *$\mathcal{D}_P$ has the properties 1, 2, and $3'$ above. Moreover,*

1. *The bisector of every pair of points consists of disjoint simple curves.*
2. *Curves portions of different bisectors intersect a finite number of times.*

We next follow the approach of [KW] and show additional properties of a polygon-offset distance function and the respective nearest-site Voronoi diagram.

Theorem 4. *Every cell of $\mathcal{V}_P^n$ is connected and $|\mathcal{V}_P^n| = O(nm)$.*

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 5. *$\mathcal{D}_P$ is complete and for every pair of points p, q in the plane there exists a point $r \notin \overline{pq}$ 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 whole plane with $O_{P,\varepsilon}$ by ranging ε from $-\varepsilon_0$ 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 whole segment $\overline{pq}$ is dominated by some edge $e_i \in P$ (where the normal to e_i and $\overrightarrow{pq}$ form the angle θ_i), then all the points $r \in \overline{pq}$ 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 $\overline{pq}$ is dominated by more than one edge of P. Assume that $\overline{pq}$ 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 2.2 that $\cos(\theta_i) < \cos(\theta_j)$ and that there exists a point $s \in \overline{pq}$ near q such that $\mathcal{D}_P(p,s) + D_P(s,q) < D_P(p,q)$. It is fairly easy to find another point s' in the plane for which $\mathcal{D}_P(p,s') + D_P(s',q) > D_P(p,q)$: select, 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. According to the mean-value theorem there exists a point r in the plane for which $f(r) = \mathcal{D}_P(p,q)$ and the claim follows. □

Note that the theorem does not require r to belong to $\overline{pq}$. The next theorem establishes a crucial extension property. (Its proof is similar to that of Theorem 5 and omitted in this extended abstract.)

Theorem 6. *For every pair of points p, q in the plane there exists a point $r \neq q$ such that $\mathcal{D}_P(p,q) + \mathcal{D}_P(q,r) = \mathcal{D}_P(p,r)$.*

Theorem 7. *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). $\mathcal{D}_P$ is complete, for every pair of points p, q in the plane there exists a point $r_1 \notin \overline{pq}$ such that $\mathcal{D}_P(p,r_1)+\mathcal{D}_P(r_1,q) = \mathcal{D}_P(p,q)$ (by Theorem 5) 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 6). This is sufficient for applying Theorem 4.1 of [KW] and obtaining the claim. □

3.2 Computing the Diagram

We define (following [KW]) the $(\alpha, \overline{\alpha})$-support as follows. A distance function D has an $(\alpha, \overline{\alpha})$-support (for $\alpha \neq \overline{\alpha} \in [0, \pi]$) if it fulfills the following condition: Let p, q be a pair of points in the plane, where $\overrightarrow{pq}$ is of slope α, and let ℓ be a line of slope $\overline{\alpha}$ 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 Figure 5).

Lemma 8. *There exist angles $\alpha \neq \beta$ s.t. $\mathcal{D}_P$ has both $(\alpha, \overline{\alpha})$- and $(\beta, \overline{\beta})$-support.*

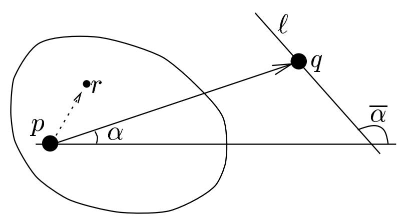

Fig. 5. $(\alpha, \overline{\alpha})$-support

Proof. There exist three edges $e_i, e_j, e_k \in P$ that never disappear in inner offsets of P (in a degenerate case mentioned earlier 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 $\overline{\alpha}$ to the slope of e_i. It is trivial to verify that $\mathcal{D}_P$ has an $(\alpha, \overline{\alpha})$-support. Similarly set β to the slope of ℓ_j and $\overline{\beta}$ to the slope of e_j. Obviously $\alpha \neq \beta$ and the claim follows. □

Now we apply the algorithmic framework of [KW] to establish the following:

Theorem 9. *Let S be a set of n points in the plane. A compact representation of $\mathcal{V}_P^n(S)$ can be computed in $O(n \log n \log m + m)$ expected time.*

Proof. The needed topological properties of $\mathcal{D}_P$ have already been established. This is sufficient for applying Theorem 5.1 of [KW] and obtaining the claim for a constant m. The extra $\log m$ factor comes from the cost of the primitive function (computing a Voronoi vertex), which is called (due to the randomized divide-and-conquer paradigm) $O(n \log n)$ times. Bisecting curves are then represented implicitly by the relative positions of the sites that define them. □

4 Compact Nearest-Site Polygon-Offset Vor. Diagram

We now extend the work of McAllister, Kirkpatrick, and Snoeyink [MKS]. They describe in detail an algorithm for constructing a *compact* nearest-site Voronoi diagram where the underlying scaled distance function is defined by a convex polygon. 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 sites), and each polygonal site is replaced by the convex hull of vertices of the site in which spokes occur. (See [MKS, p. 81, Fig. 4] 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.

To save space in this extended abstract, we assume the familiarity of the reader with [MKS] and sketch how to generalize the compact Voronoi diagram to a convex polygon-offset distance function. In order to do that we have to: (1) Make sure that the geometric properties of the compact Voronoi diagram are preserved when we change the distance function to convex polygon-offset; and (2) Verify that we can compute the same information by using the same amount of time when we make this change.

For the first goal we need only note 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 (Lemmas 3.1–3.4 and Theorems 3.5, ibid.) We omit here the details for lack of space. For the second goal we prove the following theorem:

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

Proof. We observe first how P is used for measuring regular convex-distance between points in the plane. 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 hits $\overrightarrow{pq}$. 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 should answer, given a vector $\overrightarrow{pq}$ positioned such that p coincides with the center of P, in which region of the medial-axis q falls. This region (which corresponds to a pair of vertices of P) defines the edge of P that dominates q. The elegance of this approach is in that "disappearing" edges are reflected by regions that correspond to pairs of non-consecutive vertices along P. The query thus requires $O(\log m)$ time. □

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 function `vertex`(ABC) (given three sites A, B, and C, it returns the Voronoi vertex around which they appear counter-clockwise) can still be implemented so that it runs in $O(\log n + \log m)$ time. This function is called $O(n)$ times (due to the sweep paradigm), resulting in an $O(n(\log n + \log m) + m)$-time method.

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, Meiser, and Rasch [MMR] which follows Klein's unifying approach [KW, Kl]. 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 $M_P(p,q)$ contains all the points that are closer to p than to q with respect to $\mathcal{D}_P$.

Let $S = \{p_i | 1 \le i \le n\} \in \mathbb{R}^2$ be a set of n points. The family $\mathcal{M} = \{M(p_i,p_j) | 1 \le i \neq j \le n\}$ is called a *dominant* system if for all $p_i, p_j \in S$:
1. $M(p_i,p_j)$ is open and nonempty.
2. $M(p_i,p_j) \cap M(p_j,p_i) = \emptyset$ and $\partial M(p_i,p_j) = \partial M(p_j,p_i)$.
3. $\partial M(p_i,p_j)$ is homeomorphic to the open interval $(0,1)$.

Theorem 11. *$\mathcal{M}_P$ (the family $\mathcal{M}$ with respect to $\mathcal{D}_P$) is a dominant system.*

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

A dominance system is called *admissible* if it satisfies in addition:
4. Bisectors intersect finitely-many time.
5. For all $S' \subset S, S' \neq \emptyset$ and for all reordering of indices of points in S:
   (a) Every Voronoi cell is connected and has a nonempty interior.
   (b) The union of all the Voronoi cells is the entire plane.

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

Theorem 12. *$\mathcal{M}_P$ is admissible.*

The proof of Theorem 12 is straightforward (part of the theorem was proven in a previous section) and is omitted in this version of the paper.

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

Theorem 13. *[MMR, Lemma 1] If $\mathcal{M}_P$ is semi-admissible then so is $\mathcal{M}_P^*$. Moreover, the (so-called furthest-site) Voronoi diagram that corresponds to $\mathcal{M}_P^*$ is identical to the (nearest-site) Voronoi diagram that corresponds to $\mathcal{M}_P$.*

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 [Kl, KW, 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 $O(n \log nK)$ time, where K is the time needed by a primitive function which considers 5 sites. Since computing the intersection between two bisectors requires $O(\log m)$ time, this primitive requires also this amount of time in our setting. Hence the overall expected time required for computing the furthest-site Voronoi diagram (with respect to $\mathcal{D}_P$) is $O(n \log n \log m + m)$.

We have followed the approach of [KMM] only for proving the properties of the Furthest-site Voronoi diagram. Our real goal is to obtain an optimal deterministic algorithm for computing the diagram. For this purpose we adopt the plane-sweep (in 3-space) approach of Rappaport [Ra] which follows Fortune's algorithm [Fo], in which 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 whose intersection can be computed in $O(\log m)$ time. Thus we spend $O(n \log n + m)$ time in preprocessing and $O(n(\log n + \log m))$ time in the sweep (n events, for each we spend $O(\log m)$ time for computing and $O(\log n)$ time for queue operations), resulting in an optimal $\Theta(n(\log n + \log m) + m)$-time algorithm. We again store the diagram compactly, representing bisectors implicitly via the relative positions of the two sites that define them.

6 Conclusion

We develop in this paper the notion of a polygon-offset distance function. This is an extremely important notion for tolerancing in manufacturing. We note not only that this is not a metric, it does not even fulfill the triangle inequality. 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.

Acknowledgements. The authors wish to thank J. Snoeyink and J. Bose for helpful discussions on Voronoi diagrams. We also thank V. Mirelli for introducing us the concept of polygon offsets and their applications.

References

[AA] O. Aichholzer and F. Aurenhammer, Straight skeletons for general polygonal figures in the plane, *Proc. 2nd COCOON*, 1996, 117-126, *LNCS 1090*, Springer Verlag.

[AAAG] O. Aichholzer, D. Alberts, F. Aurenhammer, and B. Gärtner, A novel type of skeleton for polygons, *J. 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, these proceedings, 1997.

[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; Short version: *Proc. 1st Symp. on Comp. Geom.*, 1985, 324–244.

[Fo] S. Fortune, A sweepline algorithm for Voronoi diagrams, *Algorithmica*, 2 (1987), 153–174.

[KN] J.L. Kelley and I. Namioka, Linear Topological Spaces, Springer Verlag, 1976.

[KS] D. Kirkpatrick and J. Snoeyink, Tentative prune-and-search for computing fixed-points with applications to geometric computation, *Fund. Informaticæ*, 22 (1995), 353–370.

[Kl] R. Klein, Concrete and abstract Voronoi diagrams, *LNCS 400*, Springer Verlag, 1989.

[KMM] R. Klein, K. Mehlhorn, and S. Meiser, Randomized incremental construction of abstract Voronoi diagrams, *Comp. Geometry: Theory and Applications*, 3 (1993), 157–184.

[KW] R. Klein and D. Wood, Voronoi diagrams based on general metrics in the plane, *Proc. 5th Symp. on Theoret. Aspects in Comp. Sci.*, 1988, 281–291, *LNCS 294*, Springer Verlag.

[MKS] M. McAllister, D. Kirkpatrick, and J. Snoeyink, A compact piecewise-linear Voronoi diagram for convex sites in the plane, *D&CG*, 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, Germany.

[Ra] D. Rappaport, Computing the furthest site Voronoi diagram for a set of disks, *Proc. 1st Workshop on Algorithms and Data Structures*, *LNCS*, 382, Springer Verlag, 1989, 57-66.