The Voronoi Diagram of Curved Objects*

Helmut Alt
Freie Universität Berlin, Institut für Informatik
Takustr. 9, D-14195 Berlin

Otfried Schwarzkopf
Universiteit Utrecht, Vakgroep Informatica
Postbus 80.089, NL-3508 TB Utrecht

Abstract

Voronoi diagrams of curved objects can show certain phenomena that are often considered artifacts: The Voronoi diagram is not connected; there are pairs of objects whose bisector is a closed curve or even a two-dimensional object; there are Voronoi edges between different parts of the same site, (so-called self-Voronoi-edges*); these self-Voronoi-edges may end at seemingly arbitrary points, and, in the case of a circular site, even degenerate to a single isolated point.*

We give a systematic study of these phenomena, characterizing their differential geometric and topological properties. We show how a given set of curves can be refined such that the resulting curves define a "well-behaved" Voronoi diagram. We also give a randomized incremental algorithm to compute this diagram. The expected running time of this algorithm is $O(n \log n)$.

1 Introduction

Voronoi diagrams are among the most extensively studied objects in computational geometry (see for instance Aurenhammer's survey [1] or the book by Okabe, Boots, and Sugihara [13]). Naturally the first type of Voronoi diagrams being considered was the one for point sites and the Euclidean metric in two dimensions. Subsequent research was concerned with generalizations of all of these features. In the two-dimensional case these generalizations were particularly motivated by applications in motion planning which lead to the so-called *retraction method* [12].

This method makes use of the fact that if there is a collision-free motion of a disk-shaped object within a collection of obstacles, from a source to a target position then there is also one that essentially follows the edges of the Euclidean Voronoi diagram of the obstacles. Since in general the obstacles are not single points, Voronoi diagrams for other types of sites were investigated, mostly for line segments. Also, if the object to be moved is not a disk but some other convex body B the retraction approach can be applied to translational motions [9]. In this case the Euclidean distance has to be generalized to a *convex distance function* that has B, with some fixed reference point inside, as its "unit circle".

The Voronoi diagram under this distance function is usually referred to as the *B-Voronoi diagram.* Klein [5, 6] gave a unified approach for many of the different variants of two-dimensional Voronoi diagrams, the so-called *abstract Voronoi diagrams.* They are not specified by distance functions but by certain topological conditions which the vertices and edges have to satisfy. Klein, Mehlhorn, and Meiser [7] gave a general paradigm for a randomized $O(n \log n)$ algorithm for constructing an abstract Voronoi diagram for a set of n sites.

Concerning the construction of B-Voronoi diagrams, so far it was mostly assumed that the obstacles are bounded by line segments and B is a convex polygon. In fact, more complex shapes can be approximated by polygons to arbitrary precision. However, in general a good approximation requires very many line segments leading to large running times of the construction algorithms. Therefore it should be interesting to consider the construction of B-Voronoi diagrams where the sites and B are bounded by more general curves. Yap [17] solves the problem for the Euclidean metric and second degree curves. Further steps in a more general direction

*This research was supported by the Netherlands' Organization for Scientific Research (NWO), by the ESPRIT Basic Research Action No. 7141, Project ALCOM II, and partially by the ESPRIT III Basic Research Action No.6546 (PROMotion). Part of this research was done while O.S. visited Freie Universität Berlin.

are made in [18], where an idea of an algorithm is given for the case that the bounding curves are circular arcs or line segments.

Unfortunately, Voronoi diagrams of curved objects do not satisfy the conditions of abstract Voronoi diagrams. Figure 1 shows the particularities that can occur even in the Euclidean case. Here we simply define the Voronoi

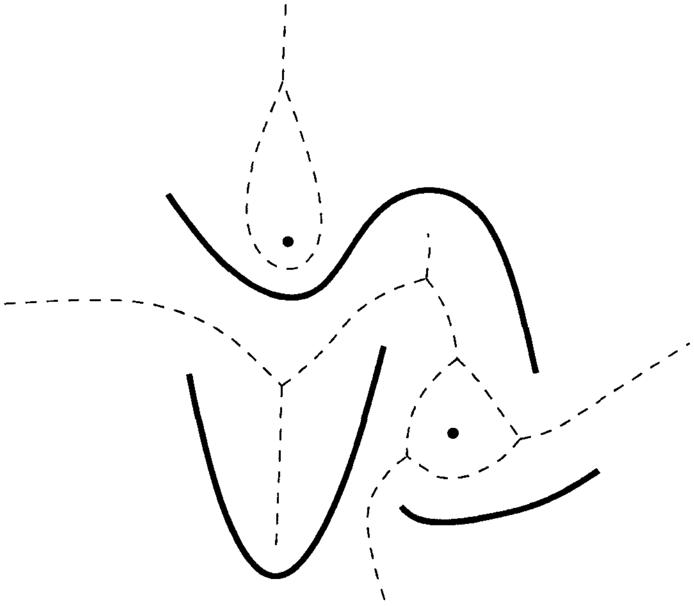

Figure 1: A Voronoi diagram of curves

diagram as the set of all points having more than one closest point on the union of all sites: The Voronoi diagram is not connected and there are Voronoi edges between different parts of the same site, which we will call *self-Voronoi-edges*; these self-Voronoi-edges may end at seemingly arbitrary points, and, in the case of a circular site, even degenerate to a single isolated point. Furthermore the bisector between two objects may be a closed curve. In the case of convex distance functions additionally there may be even pairs of objects whose bisector is a two-dimensional region.

The aim of this paper is twofold:

1. To investigate all the mentioned phenomena of Voronoi diagrams of curves and characterize their differential geometric and topological properties.

2. To show how these difficulties can be overcome (under certain preconditions) by breaking up the curves into so-called "harmless" pieces so that the idea of randomized incremental construction can be applied.

The algorithmic result of this paper is an efficient randomized algorithm for constructing the B-Voronoi diagram of a set of curves. B is assumed to be a convex body bounded by finitely many harmless curves. The algorithm will consider curves as abstract objects and assume that certain elementary operations are available as black boxes. These include finding the points having the same distance from three given sites, finding all points of a given slope, finding points where the curvature has a local maximum, given the representations of two curves finding the representation of a bisector, and finding intersection points of given curves. The details of these operations including numerical problems involved will depend on the particular application of our paradigm. For example, if it is applied to algebraic curves of some fixed degree, the elementary operations would consist of solving systems of algebraic equations of constant degree and constantly many variables.

As far as we know, this is the first systematic treatment of the phenomenon of self-Voronoi-edges. Self-Voronoi-edges may long have been considered as artifacts. We argue that self-Voronoi-edges play an essential role in Voronoi diagrams of curved objects. If the Voronoi diagram is used to do motion planning using the retraction method, for instance, then the self-Voronoi edges are necessary to capture the connectivity of the workspace. Without them, the robot may not be able to reach concavities formed by a single curve.

In this extended abstract, we can show our results in more detail for the case of the Euclidean metric only. We will give an idea as to how our results can be generalized to convex distance functions.

We will first characterize sets of curved sites that induce "well-behaved" Voronoi diagrams. As the conditions we will impose cannot be expected to hold for a given set of curves that may arise in an application, we then describe how such a given set of curves can—under some mild conditions—be refined by cutting up the curves and adding point sites on curves. We will then describe a randomized incremental algorithm of running time $O(n \log n)$ to compute the Voronoi diagram of a set of "harmless" curves. Combined with our technique to refine curves that are not yet sufficiently well-behaved, this will give us an algorithm to compute the Voronoi diagram of a set of arbitrary curves in time $O(n \log n)$.

2 The Voronoi diagram of harmless curves

A *curve* is given by a function $\gamma : I \to \mathbf{R}^2$ where $I \subset \mathbf{R}$ is some closed interval. Unless stated otherwise we will assume that curves are *regular* in the differential geometric sense, that is γ is twice continuously differentiable and $\gamma'(t) \neq 0$ for all $t \in I$. We say that two curves *touch* each other in some point $p \in \mathbf{R}^2$ iff they both pass through p without properly intersecting there. The radius of *curvature* at some point $\gamma(t)$ is (informally) the radius of the largest circle touching (and not intersecting) γ at $\gamma(t)$ on the "concave side" of γ; it is positive if that circle lies left of the direction of γ given by the parametrization and negative otherwise. The *curvature* $\gamma(t)$ is the reciprocal of the radius of curvature (for more details see Stoker [16]).

We will assume that curves are *simple*, that is $\gamma(t) \neq \gamma(t')$ for $t \neq t'$. Furthermore, each curve will not be considered as one, but as three sites: the two endpoints and the interior of the curve. We call an open curve (a curve without its endpoints) *harmless* if there is no circle that touches it in more than one point. A *harmless site* is either a point, an open circular arc, or a harmless curve. A *harmless site collection* is a finite set S of pairwise disjoint harmless sites with the condition that for every circular arc and harmless curve $\gamma \in S$ its endpoints are also members of S. For example, the parabolic arc $\{(t, t^2) \mid -3 < t < 3\}$ is not a harmless curve. We can cut it as its apex by adding the point site $(0, 0)$, and can obtain a harmless site collection of two parabolic arcs and three point sites. Observe that it is possible that several curves share one endpoint, so we allow arbitrary planar subdivisions by regular curves. Curves may not intersect but this case can be handled by making the intersection points additional point sites.

Throughout this paper, we will denote by $d(x, y)$ the Euclidean distance of points $x, y \in \mathbf{R}^2$ and for $A \subset \mathbf{R}^2$, $x \in \mathbf{R}^2$ we define $d(x, A) := \inf_{y \in A} d(x, y)$. Also, let $\psi_A(x)$ be that point of A with $d(x, \psi_A(x)) = d(x, A)$. ψ_A is not defined whenever there is no such point or when there is more than one. The following lemma is concerned with the set of points closest to some given point on a harmless curve.

Lemma 1 *Let γ be an open circular arc or a harmless curve and let p, q be its endpoints. Then:*

a) ψ_γ is defined for all x where $d(x, p) > d(x, \gamma)$ and $d(x, q) > d(x, \gamma)$, and ψ_γ is continuous wherever it is defined.

b) Let a be a point on γ, c the center of the circle of curvature at a. Only points on the ray $\overrightarrow{cp}$ are mapped onto a by ψ_γ. Other points on the same straight line, which is the normal of γ in a, that are not mapped to a are closer to one of the endpoints than to γ.

Proof: a) By definition, γ, together with its endpoints, is a compact subset of $\mathbf{R}^2$. Consequently, by continuity $a \mapsto d(x, a)$ assumes its minimum on $\gamma \cup \{p, q\}$ for some $a_0 \in \gamma$. By the harmlessness of γ there is only one such a_0, so ψ_γ is defined. Let now $x \in \mathbf{R}^2$ be some point where $a := \psi_\gamma(x)$ is defined and let $r := d(x, p)$. For some arbitrary $\varepsilon > 0$ consider the ε-neighborhood $U_\varepsilon(a)$ and the rest of the curve $A := \gamma \setminus U_\varepsilon(a)$. Let ε be sufficiently small so that $A \neq \emptyset$. Since for any $b \in A$, $d(x, b) > r$ and A is compact, there is some $\delta > 0$ such that

$$d(x, b) > r + 2\delta \tag{1}$$

for all $b \in A$. Let y be any point in $U_\delta(x)$ so that $\psi_\gamma(y)$ is defined. Since by the triangle inequality $d(y, \gamma) \leq r + \delta$, and by (1) $\delta(y, b) > r + \delta$ for any $b \in A$, it must be $\psi_\gamma(y) \in U_\varepsilon(a)$. This shows the continuity of ψ_γ.

b) Consider $x \in \mathbf{R}^2$ with $\psi_s(x) = p$. The circle around x with radius $d(x, p)$ touches s in p. Therefore the line segment $\overline{xp}$ is part of the normal through p and x does not lie beyond c, since then the circle would contain parts of s in its interior. □

Figure 2 shows the partition of the plane into 3 regions depending on whether the closest point is p, q or in the interior of γ. By Lemma 1 b) the Voronoi region of γ cannot go beyond the curve η which is the locus of all centers of curvature, the so-called *evolute* of γ.

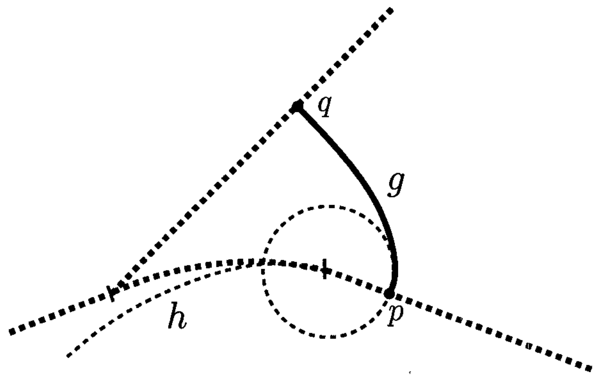

Figure 2: The Voronoi regions of p, q, and γ.

Given a harmless site collection S and a site $s \in S$, we define the *Voronoi region* of s in S, $VR(s, S)$ as follows:

$$VR(s, S) := \{x \in \mathbf{R}^2 \mid \exists p \in s \text{ with } d(x, p) = \min_{s' \in S} d(x, s')\}.$$

Note the slight twist in this definition. Had we chosen the simpler condition "$d(x, s) = \min_{s' \in S} d(x, s')$", then the Voronoi region of a curve would contain the Voronoi region of its endpoints. Our definition avoids that.

We will find it helpful to restrict our attention to the "finite" part of the Voronoi diagram. Therefore, we will add a large circle ω with center in the origin to our site collection. The radius of this circle will be assumed so large that it contains all *Voronoi vertices* and furthermore the circle is not their nearest site. The Voronoi vertices are the points whose shortest distance to any site is assumed to at least three different sites.[1] This means that nothing from the topological structure of the Voronoi diagram gets lost by inserting ω. So we

[1] Such a radius exists, in fact, it can be shown by an easy geometric argument that any triple of sites has at most two such points.

only have to consider the Voronoi diagram inside ω and, consequently, all Voronoi regions are bounded.

The *Voronoi diagram* $V(S)$ is defined as the set of all Voronoi regions $VR(s, S)$, for $s \in S$. Next we investigate the shape of Voronoi regions:

Lemma 2 *Let S be a harmless site collection, $s \in S$.*

a) If s is a point then the intersection of any straight line through s with $VR(s, S)$ is a line segment.

b) If s is a curve and p a point on S, then the normal through p intersects $\partial VR(s, S)$ in exactly two points $h_l(p)$ and $h_r(p)$ lying on either side of p. The line segment ℓ in between belongs completely to $VR(s, S)$, in fact, these are exactly the points having p as their closest point on any site. The interior points of ℓ do not belong to any other Voronoi region.

c) h_l and $h_r : s \to \mathbf{R}^2$ are continuous.

Proof: [Proof of Part b), a) follows by the same argument]
Consider some intersection point q of the normal n with $\partial VR(s, S)$. Then the circle C through p with center q contains no point of any site in its interior but another one $r \neq p$ on its boundary. Suppose there is some point $t \in n$ beyond q (seen from p) that belongs to $\partial VR(s, S)$. Then by Lemma 1b) p must be the point on s closest to t as well. But the circle with center t through p contains r, a contradiction. On the other hand, for any point $u \in \overline{pq}$ the circle around u through p is inside C, so it is empty, therefore $u \in VR(s, S)$.

c) Suppose that h_l is not continuous. So there exists a sequence of points in s converging to p whose h_l-values do not converge to $h_l(p)$. Since the Voronoi region of s is compact there exists a subsequence converging to some $q \neq h_l(p)$. Because of the continuity of ψ_s q lies on the straight line through p and $h_l(p)$. Since ∂V is closed $q \in \partial V$. This is a contradiction to part b). ◻

We summarize important topological properties of Voronoi diagrams in the following theorem:

Theorem 3 *Let S be a harmless site collection of n sites.*

(i) The union of the Voronoi regions covers ω, and no Voronoi region is empty.

(ii) For $R \subset S$ and $s \in R$ we have $VR(s, S) \subseteq VR(s, R)$.

(iii) The intersection of two Voronoi regions lies on the boundary of both.

(iv) A Voronoi region $VR(s, S)$ is simply connected, for point sites it is even star-shaped.

(v) The boundary of each Voronoi-region $VR(s, S)$ is a Jordan-curve except if s is an endpoint where several curves meet. In this case the Voronoi-region might be a line segment or the point itself.

Proof: (i) Since the union of all sites is a compact set, for any point $x \in \omega$ the minimum distance to that union is assumed, so x lies in some Voronoi region. Any Voronoi region contains the site itself and, therefore is not empty.

(ii) Let $s \in R \subset S$, and let $x \in VR(s, S)$. This means that there is a point $p \in s$ with $d(x, p) \leq d(x, s)$ for all $x \in S$. Clearly this implies $x \in VR(s, R)$.

(iii) Let $x \in VR(s, S) \cap VR(t, S)$. So there must be points $p \in s$, $q \in t$ lying on the maximal empty circle around x. By Lemma 2b) any point in the interior of the line segments $\overline{xq}$ and $\overline{xp}$ lies in $VR(s, S) \setminus VR(t, S)$ and $VR(t, S) \setminus VR(s, S)$, respectively. This shows that x lies on the boundary of both.

(iv) Let α be some closed curve within $VR(s, S)$. Let q be some arbitrary point surrounded by α and n the normal to s through q which is unique by Lemma 1b) (the straight line sq if s is a point). n intersects α in at least two points $a, b \in VR(s, S)$ surrounding q. Then by Lemma 2b) $q \in VR(s, S)$. Since this holds for any point q in the region encircled by α, α is within $VR(s, S)$ contractible to a point. Since this holds for any α, $VR(s, S)$ is simply connected.

(v) Let s be a harmless curve and consider it oriented by its parametrization. Then to any $x \in S$ there exist unique points $h_\ell(x)$, $h_r(x) \in \partial VR(x, S)$ to the left and right of s, respectively with $\psi_s(h_\ell(x)) = \psi_s(h_r(x)) = x$ (see Lemma 2b)). Since h_ℓ and h_r are continuous (Lemma 2c)) and because of the harmlessness of s, they are also one-to-one the images $A := h_\ell(s), B := h_r(s)$ are homeomorphic to s, so they are simple curves. By continuity reasons their endpoints have the endpoints p, q of s as closest points. Altogether, the boundary of $\partial VR(s, S)$ consists of A, B and two segments I, J (that may degenerate to a single point) of the normals through p, q, respectively (see Figure 3). Since I, A, J, B are pairwise non-intersecting, their

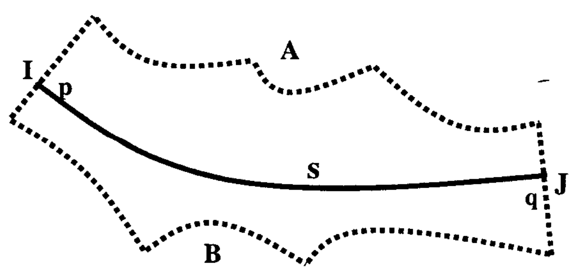

Figure 3: Voronoi region $VR(s, S)$

concatenation forms a Jordan curve. If s is a circular segment the argument is similar only that one of A, B may degenerate to a single point.

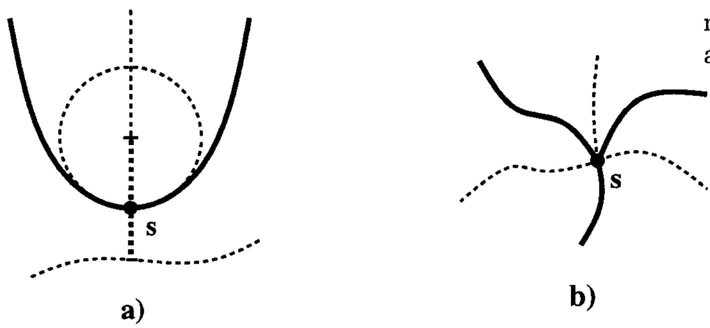

Figure 4: Degenerate Voronoi-cells

Similar techniques prove the statement, if s is an isolated point or an endpoint of a curve. If several curves share s as an endpoint $VR(s, S)$ may degenerate to a line segment (Figure 4a)) or just s itself (Figure 4b)). The former phenomenon occurs when the closed half-spaces beyond the normals of those curves in s intersect only in a straight line, the latter when they intersect only in s itself. □

The Voronoi diagram can also be represented as a graph as follows: The vertices of the graphs of are the points of $V(S)$ which belong to the boundary of three or more Voronoi regions. As was mentioned before, it can be shown that there are at most two such points per triple of sites. The edges of the graph correspond to the maximal connected subsets belonging to the boundary of exactly two Voronoi regions. They are curves by Theorem 3(v). The faces of the graph correspond to the Voronoi regions. We will use $V(S)$ to denote the graph as well. Using Theorem 3(iv) and Euler's formula we can prove:

Theorem 4 *Given a harmless site collection S of n sites. The Voronoi graph $V(S)$ is a planar connected graph with at most $n+1$ faces, at most $3n-3$ edges, and at most $2n-2$ vertices.*

Here the outer circle ω is not counted as a site, but the edges and vertices where it is involved are counted.

3 Partitioning curves into harmless pieces

Here we will see that the points responsible for self-Voronoi-edges or, in other words, the non-harmlessness of curves are the points where the absolute value of the curvature has a local maximum. Figure 5 shows this situation. There is a circle around a touching γ in two points. When decreasing the radius the center of the circle traces a self-Voronoi-edge that ends at c where the two tangent points fall together. c is the center of the circle of curvature C of γ at p which is a local maximum of curvature. Points on the line segment $\overline{cp}$ are centers of circles that touch γ in p.

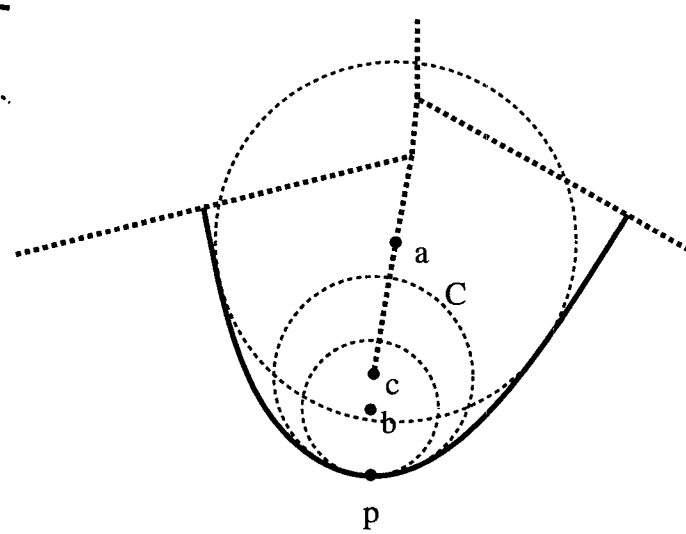

Figure 5: A local maximum of curvature

If we remove these local maxima by cutting the curve there and making them separate sites, the Voronoi diagram will have the nice properties described in Theorem 3. In fact, it holds

Theorem 5 *A regular curve that does not contain its endpoints is harmless if it has no two parallel tangents, it contains no circular segments, and the absolute value of its curvature has no local maximum.*

Proof: Let γ be such a curve. Since it has no two parallel tangents we can assume wlog. that it has no vertical tangents. So it can be parametrized as the graph of a function $f : I \to \mathbf{R}$, where $I \subset \mathbf{R}$ is some interval, and $\gamma(t) = (t, f(t))$. Suppose some circle C is not intersecting γ but touching it in two points p and q. Since there is no vertical tangent C lies in p and q on the same side of γ. Without loss of generality we can assume that this is the left side, so the curvature of γ is nonnegative, and that both p and q lie in the lower semi-circle of C. Now the curvature of γ is at most the curvature of C in p and q, the latter is constant, and the former has no maximum in between. It follows that the curvature of γ is not larger than the curvature of C between p and q. Because of the following lemma this is only possible if γ coincides with C between p and q, a contradiction. □

Lemma 6 *Let γ, δ be two regular curves that are graphs of functions $f, g : I \to \mathbf{R}$ where $I \subset \mathbf{R}$ is some* *finite interval, so* $\gamma(t) = \big(t, f(t)\big)$, $\delta(t) = \big(t, g(t)\big)$ *for all* $t \in I$. *If* γ *touches* δ *in two points* $t_1, t_2 \in I$ *and its curvature is not greater than* δ*'s for all* $t \in [t_1, t_2]$ *then* γ *and* δ *must coincide in* $[t_1, t_2]$.

This lemma follows from the fact, that the curvature of the graph of a function f is given by

$$\kappa(t) = \frac{f''(t)}{\left(1 + \left(f'(t)\right)^2\right)^{3/2}}$$

by standard analytical considerations.

In Section 5 we will see how to compute the Voronoi diagram of a harmless site collection in time $O(n \log n)$. Theorem 5 allows us to apply that algorithm to more general sets of curves: In fact, it allows us to partition the given set of curved objects into harmless pieces and circular arcs. If, for instance, the given curves are algebraic of constant degree, each can have at most a constant number of points of vertical tangency or maxima of the curvature. By cutting the n original curves at these points we obtain a collection of $O(n)$ harmless curves. We can then compute the Voronoi diagram of these harmless pieces, and obtain a Voronoi diagram of complexity $O(n)$. If that is desired, we can then merge the Voronoi cells of curves that are pieces of the same original curve. In most applications, however, that is probably not what is needed: If the Voronoi diagram is used for motion planning with the retraction method, for instance, the additional self-Voronoi-edges are essential to guarantee that the resulting road map captures the connectivity of the workspace.

4 Convex distance functions

Let $B \subset \mathbf{R}^2$ be some convex body with some fixed reference point o inside. The *B-distance* from a point $p \in \mathbf{R}^2$ to a point q is defined as the factor by which B, when placed at p with o, has to be stretched around o until it touches q. The B-distance is in general not symmetric, but it satisfies the triangle inequality.

Here we will assume that B is bounded by a constant number k of harmless curve segments. So the boundary of B consists of $2k$ *features*: The k harmless curves and k vertices where they meet. We further partition the curve sites into pieces so that one piece can be touched only by one feature of B. This can be done as follows: To each feature of B there is an interval of possible tangent slopes. We partition the sites at points where the slope equals one of the interval boundaries (see Figure 6). Notice that we have to distinguish between the two sides of a curve and partition both accordingly. The number of sites produced this way equals the original number of sites plus the number of places on sites where one of the tangent slopes occurs. This number cannot be bounded in general, but for example for algebraic curves of constant degree the increase is by a constant factor.

In order to derive our results from the previous section to arbitrary B-distance we have to generalize our definitions: The *B-curvature* of a regular curve γ in a point p is the largest copy of B that touches γ in p without intersecting it. If the feature of B touching p is a curve δ, the B-curvature is the ratio of the curvatures of γ and δ at that point. If the feature is a non-smooth vertex, the B-curvature is 0, no matter what γ looks like.

A *B-harmless* site is a curve that cannot be touched by any homothet of B in more than one place. A *B-harmless site collection* consists of open B-harmless sites, their endpoints, other points, and curves that are homothets of segments of the boundary of B.

The *B-normal* of a curve γ in a point p is the orbit of the reference point o considering all homothets of B that touch γ in p (see Figure 7). Any B-normal consists of a ray on either side of γ.

With these definitions, all results and proofs of the previous section go through, if the boundary of B contains no straight segments.

Otherwise, the following problems will occur: Let e be a straight line segment on the boundary of B. Then for any point p on some site where the tangent is parallel to e the B-curvature is not defined since an arbitrarily large copy of B can be placed there. (With the exception that p is an inflection point of γ, i.e. has curvature 0.) p is a site by itself, as was mentioned before. There could be more point sites all lying on a straight line parallel to e. In this case there are "two-dimensional Voronoi-edges" and vertices, i.e. two-dimensional areas of points with the same distance to several of the sites. The exact partition may have quadratic complexity.

However, as we shall see, our algorithm will first construct the Voronoi diagram of the point sites and then insert the curves. For point sites however, there is an efficient algorithm by Klein and Wood [8] manageing the mentioned degeneracies which can be applied. For the curve sites these situations cause no problems any more and the properties of Theorem 3 hold.

5 Randomized incremental construction

The Voronoi diagram of curved objects does not fit into the framework of abstract Voronoi diagrams by Klein et al. [7] and cannot be computed with their randomized incremental algorithm. The reason for this is that they assume that the bisector of any pair of sites is an unbounded simple curve. Even for a point and a circular arc, this is no longer true under the Euclidean metric: The bisector can be a closed curve.

If, however, the objects form a harmless site collec-

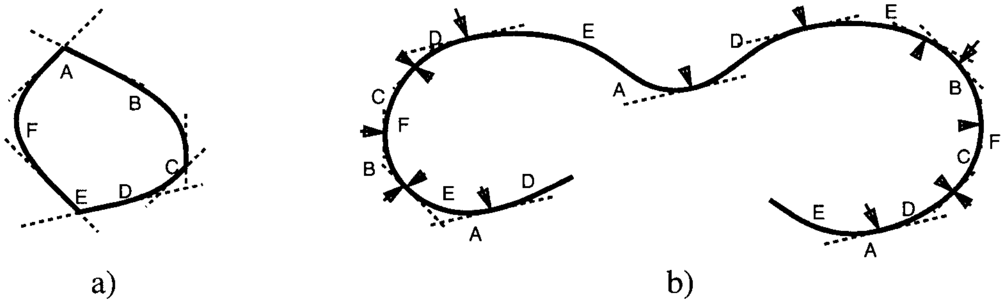

Figure 6: a) convex body B with its features b) partitioning of a curve γ

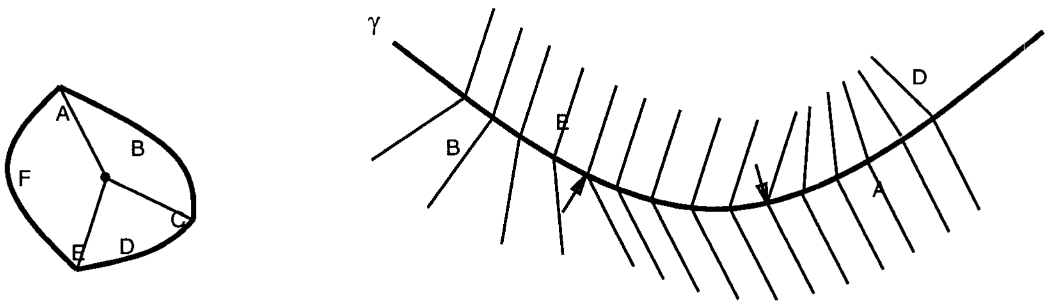

Figure 7: B-normals at γ

tion, this situation can be remedied. In the following we demonstrate how to compute the Voronoi diagram of a harmless site collection in time $O(n \log n)$ using a kind of randomized incremental algorithm, based on the framework set by Clarkson and Shor [4], Mulmuley [11], and Boissonnat et al. [2]. We have to make sure that during the algorithm which constructs the Voronoi diagram by inserting the sites one by one, the intermediate set of sites is always a harmless site collection, that is that no curve is inserted before both of its endpoints are.

Therefore, we compute the Voronoi diagram of a harmless site collection S in two stages. In the first stage, we compute a Voronoi diagram of n points $V(P \cup Q)$. Here $P \subset S$ is the set of all *point sites.* Q is obtained by selecting for each curve site $s \in S$ a point q_s in the relative interior of s.

The points in Q serve as "placeholders" for the curves they stem from which, in the second stage are added one by one in random order. This replacement of the placeholders by the actual curves is made easy by the fact that we already know where to insert a new site s. We will restrict ourselves to the description of the Euclidean case.

We will need to represent the Voronoi diagram $V(R)$ of a subset $R \subset S \cup Q$. This can be done using any standard structure for planar subdivisions, such as the doubly-connected edge list [10, 14]. In the first stage of the algorithm, we simply compute the Voronoi diagram $V(P \cup Q)$. This can be done using any efficient algorithm for the construction of Voronoi diagrams and takes time $O(n \log n)$.

As discussed in Section 2, we add a large circle ω to our site collection S. If necessary, this circle has to be handled symbolically. After computing $V(P \cup Q)$, we add ω to obtain $V(P \cup Q \cup \{\omega\})$. This can be done in time $O(n)$.

Let now $s_1, s_2, \ldots, s_m$, be a random permutation of the curve sites in S. Let q_r denote q_{s_r}, the point in the interior of s_r we had chosen. We consider the sites s_r in this order. In every step of the algorithm, we have to replace a point site q_r in the current Voronoi diagram $V(P \cup \{\omega, s_1, \ldots, s_{r-1}, q_r, \ldots, q_m\})$ by a curve site s_r to obtain $V(P \cup \{\omega, s_1, \ldots, s_r, q_{r+1}, \ldots, q_m\})$. Let's look at this in more detail.

For brevity, let $s := s_r$, let $q := q_s = q_r$, let $R := P \cup \{s_1, \ldots, s_{r-1}, q_r, q_{r+1}, \ldots, q_m\}$, and let $R' := P \cup \{s_1, \ldots, s_{r-1}, s_r, q_{r+1}, \ldots, q_m\}$. By Theorem 3, the Voronoi region $VR(s, R')$ of s is a simply connected region whose boundary is a closed Jordan-curve (recall that s is not a point site). To obtain $V(R')$ from $V(R)$ means to remove the portion $\mathcal{I}$ of $V(R)$ that lies in $VR(s, R')$, and to add the boundary of $VR(s, R')$. We first prove a lemma.

Lemma 7 *The "skeleton" $\mathcal{I}$ contains the boundary of $VR(q, R)$, is connected, and contains no cycle except for the boundary of $VR(q, R)$. All its leaves lie on the boundary of $VR(s, R')$, and its complexity is linear in* *the number of these leaves.*

Proof: Since q lies on s, clearly $VR(q,R)$ is completely contained in $VR(s,R')$, and hence $\mathcal{I}$ contains the boundary of $VR(q,R)$. If $\mathcal{I}$ contained any other cycle, some Voronoi region $VR(r,R)$, $r \in R$, $r \neq q$, must lie inside $VR(s,R')$ which would imply that r lies in $VR(s,R')$, which is impossible.

Assume now that $\mathcal{I}$ is not connected. That means that there are at least two connected components $\mathcal{I}_1$ and $\mathcal{I}_2$ of $\mathcal{I}$. By Theorem 3(iv), none of these can be contained in the interior of $VR(s,R')$, they both must have some connection with $\partial VR(s,R')$. So there is a path $\gamma \subset VR(s,R') \setminus \mathcal{I}$ connecting two points x and y on the boundary of $VR(s,R')$ and separating $\mathcal{I}_1$ from $\mathcal{I}_2$. Since $\gamma \cap V(R) = \emptyset$, γ lies in the interior of some $VR(r,R)$, for $r \in R$. This implies that there must be points x' and y' arbitrarily close to x and y that lie in $VR(r,R')$. That means that x and y can be connected by a path γ' in $VR(r,R')$. The combination $\gamma \cup \gamma'$ is a closed loop in $VR(r,R)$ containing either $\mathcal{I}_1$ or $\mathcal{I}_2$, a contradiction. It follows that $\mathcal{I}$ is connected.

Leaves of $\mathcal{I}$ must clearly lie on $\partial VR(s,R')$. By removing one edge on $\partial VR(q,R)$ from $\mathcal{I}$, it becomes a tree all of whose interior vertices have degree at least three. Consequently, its complexity is linear in the number of its leaves. ◻

Consider now the boundary of $VR(s,R')$. As discussed in the proof of Theorem 3(v), it consists of two segments I and J—possibly degenerated to a point—through the endpoints p and q of s and two simple curves A and B. The curves A and B consist of a sequence of

- edges that lie in the interior of some $VR(r,R)$, $r \in R$, and are hence part of the bisector of r and s,
- crossings between edges of $V(R)$ and $\partial VR(s,R')$, and
- vertices of $V(R)$.

We first identify I, J, and the skeleton $\mathcal{I}$. This can be done using a graph search starting at any edge on the boundary of $VR(q,R)$ and takes time linear in the complexity of $\mathcal{I}$. The leaves of $\mathcal{I}$ are the vertices of $VR(s,R')$, and from that information we can then construct the curves A and B to obtain $VR(s,R')$ in time linear in the complexity of $\mathcal{I}$ and $VR(s,R')$.

It remains to bound the running time of the algorithm. As we observed, the first stage of the algorithm takes time $O(n \log n)$. Inserting curve site s_r takes time linear in the complexity of $\mathcal{I}$ and $VR(s_r,R')$ and thus by Lemma 7 linear in the complexity of the new Voronoi region $VR(s_r,R')$.

It remains to bound the expected size of $VR(s_r,R')$. We use a standard backwards-analysis argument [3, 15]: Fix R', and let s be a random curve site in R'. The total complexity of $V(R')$ is $O(n)$ by Theorem 4, and there are r possible choices for s. Consequently, the expected complexity of $VR(s,R')$ is $O(n/r)$. Summing this over all curve sites, we find that the second stage of the algorithm takes expected time $O(n \log n)$ as well.

Theorem 8 *The two-stage randomized incremental algorithm constructs the Voronoi diagram of a harmless site collection of n sites in time $O(n \log n)$.*

Acknowledgments. We would like to thank Stefan Felsner and Lorenz Wernisch for fruitful discussions and valuable hints.

References

[1] F. Aurenhammer. Voronoi diagrams: a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, 1991.

[2] J.-D. Boissonnat, O. Devillers, R. Schott, M. Teillaud, and M. Yvinec. Applications of random sampling to on-line algorithms in computational geometry. *Discrete Comput. Geom.*, 8:51–71, 1992.

[3] L. P. Chew. Building Voronoi diagrams for convex polygons in linear expected time. Technical Report PCS-TR90-147, Dept. Math. Comput. Sci., Dartmouth College, Hanover, NH, 1986.

[4] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[5] R. Klein. Abstract Voronoi diagrams and their applications. In *Computational Geometry and its Applications*, volume 333 of *Lecture Notes in Computer Science*, pages 148–157. Springer-Verlag, 1988.

[6] R. Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes in Computer Science*. Springer-Verlag, 1989.

[7] R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom. Theory Appl.*, 3(3):157–184, 1993.

[8] R. Klein and D. Wood. Voronoi Diagrams and Mixed Metrics. Proc. 5th Sympos. Theoret. Aspects Comput. Sci. (STACS), Lecture Notes in Computer Science 294, Springer-Verlag 1988, 281–291.

[9] D. Leven and M. Sharir. Planning a purely translational motion for a convex object in two-dimensional space using generalized Voronoi diagrams. *Discrete Comput. Geom.*, 2:9–31, 1987.

[10] D. E. Muller and F. P. Preparata. Finding the intersection of two convex polyhedra. *Theoret. Comput. Sci.*, 7:217–236, 1978.

[11] K. Mulmuley. A fast planar partition algorithm, I. In *Proc. 29th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 580–589, 1988.

[12] C. Ó'Dúnlaing and C. K. Yap. A "retraction" method for planning the motion of a disk. *J. Algorithms*, 6:104–111, 1985.

[13] A.Okabe, B. Boots, and K. Sughihara. *Spatial Tesselations: Concepts and Applications of Voronoi diagrams.* J. Wiley & Sons, 1992.

[14] F. P. Preparata and M. I. Shamos. *Computational Geometry: an Introduction.* Springer-Verlag, New York, NY, 1985.

[15] R. Seidel. Backwards analysis of randomized geometric algorithms. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, volume 10 of *Algorithms and Combinatorics*, pages 37–68. Springer-Verlag, 1993.

[16] J. J. Stoker. *Differential Geometry.* Wiley-Interscience, 1969.

[17] C. K. Yap. An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. *Discrete Comput. Geom.*, 2:365–393, 1987.

[18] Chee K. Yap and Helmut Alt. Motion planning in the CL-environment. In *Proc. 1st Workshop Algorithms Data Struct.*, volume 382 of *Lecture Notes in Computer Science*, pages 373–380, 1989.