

The Voronoi Diagram of Curved Objects*

Helmut Alt,¹ Otfried Cheong,² and Antoine Vigneron³

¹Institut für Informatik, Freie Universität Berlin,
Takustraße 9, 14195 Berlin, Germany
alt@inf.fu-berlin.de

²Division of Computer Science, KAIST,
373-1 Guseong-dong, Yuseong-gu,
Daejeon 305-701, Korea
otfried@cs.kaist.ac.kr

³Department of Computer Science, National University of Singapore,
3 Science Drive 2, Singapore 117543
antoine@comp.nus.edu.sg

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 not on a site, 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 [2] or the book by Okabe et al. [19]). Naturally the first type of Voronoi diagrams 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.

* This work was supported by the the Brain Korea 21 Project, The School of Information Technology, KAIST, 2005, and by the National University of Singapore under Grant R252-000-130.

In the two-dimensional case these generalizations were also motivated by applications in motion planning which lead to the so-called *retraction method* [18]. 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. More complex shapes can be approximated by polygons to arbitrary precision. However, in general a good approximation requires very many line segments and leads to large running times of the construction algorithms. Therefore it should be interesting to consider the construction of Voronoi diagrams where the sites are bounded by more general curves.

Yap [24] solves the problem for the Euclidean metric and second degree curves. Further steps in a more general direction are made by Yap and Alt [25], who gave an idea of an algorithm for the case that the bounding curves are circular arcs or line segments. Karavelas and Yvinec [12] consider the case of convex objects in the plane.

Klein [13], [14] 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 et al. [15] gave a general paradigm for a randomized $O(n \log n)$ algorithm for constructing an abstract Voronoi diagram for a set of n sites.

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




Fig. 1. A Voronoi diagram of curves.

two objects may be a closed curve. Self-Voronoi-edges, which are part of the *medial axis* of a curve, 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.

The aim of this paper is threefold:

1. To investigate 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 “harmless” pieces.
3. To give a randomized incremental algorithm to compute the Euclidean Voronoi diagram of these pieces.

A “harmless” curve is either a line segment, a circular arc, or a spiral arc. A spiral arc is a curve whose curvature monotonically increases strictly along the curve. A circular arc can be seen as the extreme case of a spiral arc where the curvature remains constant, while a line segment can be seen as the extreme case of a circular arc whose curvature is zero. Perhaps the corner stone of our results is the insight that spiral arcs are well-behaved with respect to Voronoi diagram formation.

A set of curves that arises in an application will not necessarily consist of harmless curves only. We therefore describe how such a given set of curves can—under some mild conditions—be refined into harmless pieces by cutting up the curves and adding point sites on curves.

The algorithmic result of this paper is a randomized algorithm of running time $O(n \log n)$ for constructing the Voronoi diagram of a set of n harmless curves. Combined with our technique to refine curves that are not yet sufficiently well-behaved, this results in an algorithm to compute the Voronoi diagram of a set of arbitrary curves in time $O(n \log n)$.

The algorithm considers curves as abstract objects and assumes 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. The underlying algebraic and numerical questions are quite involved and have been investigated in several publications in the computer graphics and CAD community [5], [7]–[10].

A preliminary version of this work appeared in the *Proceedings of the ACM Symposium on Computational Geometry*, 1995 [1]. Lemma 1(b) of that version is incorrect. Overcoming the incorrect claim led us to the following exposition that is based on spiral arcs.

After the first version of this work, Farouki and Ramamurthy published a different approach [8]–[10] for computing Voronoi diagrams of curved objects. They also

have a technique for splitting curves into simpler pieces that allows one to compute self-Voronoi edges, but it is by no means trivial [8], so they restrict their attention to cubics. With our approach one only needs to find extrema of curvature, which is straightforward for arbitrary algebraic curves. Besides, the only time bound given is $O(n^2)$ [10].

2. Spiral Arcs

We assume that curves are regular in the differential-geometric sense, and that they do not contain their endpoints. More precisely, a *curve* γ is a subset of \mathbb{R}^2 such that there is a function $f_\gamma: [0, 1] \rightarrow \mathbb{R}^2$ with the following properties: f_γ is twice continuously differentiable, $f'_\gamma(t) \neq 0$ for all $t \in [0, 1]$ and $\gamma = \{f_\gamma(x) \mid 0 < x < 1\}$. The *endpoints* of γ are $f_\gamma(0)$ and $f_\gamma(1)$. We say that two curves *touch* each other in some point $p \in \mathbb{R}^2$ iff they both pass through p and have the same tangent direction there. As a special case, a circle touches a curve γ in p if it contains p and its center lies on the normal to γ through p . Consider a touching circle C whose center lies on the concave side of the curve. If the radius of C is small, it will locally—in a neighborhood of p —lie on the concave side of γ . If its radius is large, the circle lies locally on the convex side of γ . The *radius of curvature* separates the two sets of radii. In general, an *osculating circle* (the circle of curvature) touches and crosses the curve in p (the only exception is in points p where the curvature is extreme). The *curvature* is the reciprocal of the radius of curvature (for more details see for instance [22]).

A curve is a *spiral arc* if it is convex (that is, the union of the curve and the segment connecting its endpoints is a closed convex curve) and if its curvature is strictly monotone. In particular, a spiral arc does not contain a point of zero curvature. The convexity requirement excludes genuine spirals (which have arbitrarily long and complex self-Voronoi edges). Our definition of a spiral arc is more restricted than definitions given elsewhere; Guggenheimer [11], for instance, requires only monotone curvature (not strictly monotone), and does not have the convexity requirement.

The following theorem is due to Kneser [11, Theorem 3-12].

Theorem 1 (Kneser's Theorem). *Any circle of curvature of a spiral arc contains every smaller circle of curvature of the arc in its interior.*

A point $p = \gamma(t_0)$ on a spiral arc γ decomposes the arc into two branches, a branch γ_1 with increasing curvature, and a branch γ_2 with decreasing curvature when traversing these branches starting from p . Kneser's theorem implies that γ_2 lies inside the osculating circle C at p , while γ_1 lies outside C . In particular, this implies that a spiral arc is *simple*, that is $\gamma(t) \neq \gamma(t')$ for $t \neq t'$.

The following theorem is due to Vogt [23], [11, Theorem 3-17], see Fig. 2.

Theorem 2 (Vogt's Theorem). *Let p and q be two points on a spiral arc, with curvature strictly increasing from p to q . Then the angle β of the tangent to the arc at q with the chord pq is larger than the angle α of the tangent at p with pq .*




Fig. 2. Vogt's theorem.

Vogt's theorem implies the following result.

Lemma 3. *Let γ be a spiral arc. There is no circle that touches γ in more than one point.*

Proof. Assume a spiral arc γ touches circle C at p and q . The angles of the tangents to C at p and q with pq are equal. So the same is true for the tangents to γ at p and q , which is impossible by Vogt's theorem. \square

3. Harmless Sites

Recall that our goal is to give a systematic treatment of self-Voronoi edges. Our first central observation is that self-Voronoi edges are closely related to local maxima of the curvature of the curves, as Fig. 3 shows. There is a circle around x 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 in the point u . The point c is the center of the circle of curvature C of γ at u , which is a local maximum of curvature. Points y on the line segment \overline{cu} are centers of circles that touch γ only in u and, thus, do not contribute to the self-Voronoi edge.

Since spiral arcs have monotone curvature, they cannot have local maxima of the curvature. Lemma 3 implies that indeed no self-Voronoi edge can exist for a single




Fig. 3. A local maximum of curvature.

spiral arc: no disk can touch a spiral arc in two places. This insight leads us to restrict our discussion of Voronoi diagrams to sites that are spiral arcs. We also allow arcs of constant curvature, that is, circular arcs and line segments.

Formally, let a *harmless site* be either a point, a line segment, a circular arc, or a spiral arc. A *harmless site collection* is a finite set S of pairwise disjoint harmless sites with the condition that for every line segment, circular arc, and spiral arc $\gamma \in S$ the endpoints of γ are also members of S .

For example, the parabolic arc $\{(t, t^2) \mid -3 < t < 3\}$ is not a spiral arc. We can cut it at 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. Closed curves (loops) can be handled by cutting them at some points, and adding those points to the site collection.

In Section 5 we will see how to compute the Voronoi diagram of a harmless site collection in time $O(n \log n)$. To apply that algorithm to more general sets of curves, we have to partition the curves into spiral arcs, straight segments, 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 and local maxima of the curvature, and a nondegenerate interval of constant curvature is only possible if the whole curve is a circular arc. By cutting the n original curves at these points we obtain a collection of $O(n)$ harmless sites. We can then compute the Voronoi diagram of these 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. The Voronoi Diagram of Harmless Sites

Throughout this paper we denote by $d(x, y)$ the Euclidean distance of points $x, y \in \mathbb{R}^2$ and for $A \subset \mathbb{R}^2$, $x \in \mathbb{R}^2$ we define $d(x, A) := \inf_{y \in A} d(x, y)$. Also, let $\psi_A(x)$ be that point of A with $d(x, \psi_A(x)) = d(x, A)$. Whenever there is no such point or when there is more than one, then ψ_A is not defined. Let Ψ_A be the region where ψ_A is defined. On Ψ_A we define an equivalence relation \equiv , where $x \equiv y$ if and only if $\psi_A(x) = \psi_A(y)$.

Figure 4 shows the region Ψ_γ and some equivalence classes for γ a line segment, a circular arc, or spiral arc. Note that the boundary of Ψ_γ (shown dotted in the figure) is in general not part of the region, except for the curved boundary in the case of a spiral arc. Observe the fibration of Ψ_γ defined by \equiv . Parts (a) and (b) of the figure require no further proof, so we consider the case of a spiral arc.

Lemma 4. *Let γ be a spiral arc with endpoints p and q , such that the curvature increases from p to q . Let u be a point on γ and let n be the line normal to γ in u . Then $\psi_\gamma(x) = u$ if and only if x is a point on n and $d(x, u) \leq d(x, q)$.*

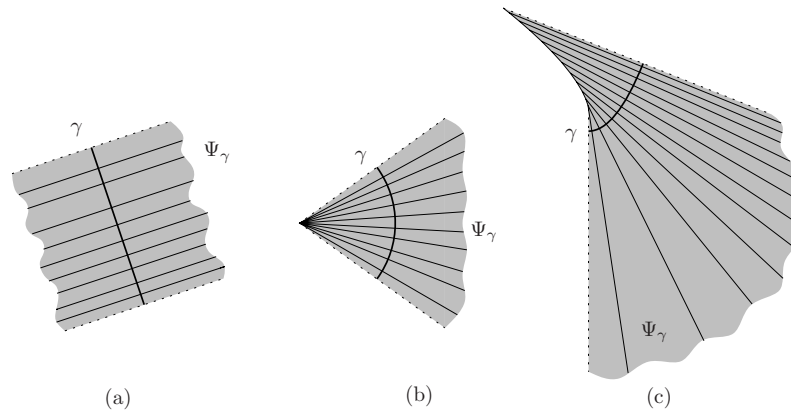


Fig. 4. The region Ψ_γ for the three different curve types.

Proof. The “only if” direction is obvious, so assume that x is a point on n with $d(x, u) \leq d(x, q)$. If x is on the convex side of γ then clearly $\psi_\gamma(x) = u$, so we also assume that x is on the concave side. We denote by C the circle through u with center x and by C_u the osculating circle at u (see Fig. 5(a)). Since both circles share a common tangent at u and q is outside the interior of C , and, by Kneser’s theorem, inside C_u , we know that the interior of C is contained in the interior of C_u . Kneser’s theorem also implies that the osculating circle at p contains C_u in its interior, so p is outside C_u and, therefore, outside C . Hence there is a closest point to x in γ , we denote this point by v . We will prove by contradiction that $v = u$, so from now on we assume $v \neq u$.

By Lemma 3, the point v cannot lie on C , so v is inside C , and thus v cannot be situated between p and u along γ . However, if v lies between u and q , there is a point w on C such that γ touches C at u from the outside, enters the interior of C , visits v , leaves C at w , and runs to q (see Fig. 5(b)). Then the angle between the tangent to γ at u and the line uw is larger than the angle between the same line and the tangent to γ at w , contradicting Vogt’s theorem. \square




Fig. 5. Proof of Lemma 4.

The lemma explains why spiral arcs are so well-behaved under Voronoi-diagram formation: to verify the nearest point on a spiral arc, we need to consider only the neighborhood of that point and the endpoint of the arc. Lemma 4 has the following corollary.

Corollary 5. *Let γ be a spiral arc with endpoints p and q , and curvature increasing from p to q . Then Ψ_γ is the region bounded by normals in p and q and by the bisector of q and γ . The latter is part of Ψ_γ , the normals are not. The center of any osculating circle of γ lies outside Ψ_γ .*

Lemma 6. *Let γ be a line segment, circular arc, or spiral arc. The function ψ_γ is continuous on Ψ_γ .*

Proof. Let $x \in \Psi_\gamma$, let $u := \psi_\gamma(x)$, and let $r := d(x, u)$. We denote by $\bar{\gamma}$ the union of γ and its endpoints. For some arbitrary $\varepsilon > 0$ consider the ε -neighborhood $U_\varepsilon(u)$ and $A := \bar{\gamma} \setminus U_\varepsilon(u)$. Let ε be sufficiently small so that $A \neq \emptyset$. Since for any $u' \in A$, $d(x, u') > r$ and A is compact, there is some $\delta > 0$ such that

$$d(x, u') > r + 2\delta \quad (1)$$

for all $u' \in A$. 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) $d(y, u') > r + \delta$ for any $u' \in A$, it must be $\psi_\gamma(y) \in U_\varepsilon(u)$. This shows the continuity of ψ_γ . \square

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:

$$VR(s, S) := \{x \in \mathbb{R}^2 \mid \text{there is a unique } p \in s \text{ with } d(x, p) = \min_{s' \in S} d(x, s')\}.$$

An example of Voronoi regions associated with a harmless site collection S is given in Fig. 6. The sites in S are a spiral arc γ , its two endpoints p and q , and another point a . In this example the edges bounding a region belong to it, except for the edges containing p and q which are not part of $VR(\gamma, S)$. Note the slight twist in the definition. Had we chosen the simpler condition

$$d(x, s) = \min_{s' \in S} d(x, s'),$$




Fig. 6. Example of Voronoi regions.

then the Voronoi region of a curve would contain the Voronoi region of its endpoints. Our definition avoids that. The uniqueness of p serves only to exclude the center of a circular arc from its Voronoi region (the point p is unique in all other cases, trivially so for line segments and circular arcs, by Lemma 3 for spiral arcs).

The *Voronoi diagram* $V(S)$ is defined as the set of all Voronoi regions $VR(s, S)$, for $s \in S$. The Voronoi diagram can also be represented as a *Voronoi graph* as follows: the vertices of the graph are the points in \mathbb{R}^2 which belong to the boundary of three or more Voronoi regions. The edges of the graph correspond to the maximal connected subsets belonging to the boundary of exactly two Voronoi regions. (In these two definitions, the boundary of a Voronoi region is defined according to the usual topology of \mathbb{R}^2 , except if the Voronoi region is a line segment, in which case we consider that its boundary consists of its two endpoints.) We show below that Voronoi edges are curves. The faces of the graph correspond to the Voronoi regions.

We will find it helpful to restrict our attention to the “finite” part of the Voronoi diagram. Therefore, we add a large circle ω with center in the origin to our site collection. The radius of this circle is assumed so large that it contains all Voronoi vertices and furthermore the circle is not their nearest site. This means that no information about the topological structure of the Voronoi diagram gets lost by inserting ω . So we only have to consider the Voronoi diagram inside ω and, consequently, all Voronoi regions are bounded.

Next we investigate the shape of Voronoi regions:

Lemma 7. *Let S be a harmless site collection and let $\gamma \in S$ be a curve. Then $VR(\gamma, S) \subseteq \Psi_\gamma$. For any $u \in \gamma$, the equivalence class $K_u := \{x \in \Psi_\gamma \mid x \equiv u\}$ of u intersects $VR(\gamma, S)$ in a straight line segment with endpoints $h_\ell(p)$ and $h_r(p)$ lying on opposite sides of u . The interior points of this segment do not belong to the closure of any other Voronoi region. The functions h_ℓ and $h_r: \gamma \rightarrow \mathbb{R}^2$ defined this way are continuous.*

Proof. If $x \in VR(\gamma, S)$, then there is $u \in \gamma$ such that $d(x, u) = \min_{s \in S} d(x, s)$. In particular, $d(x, u) = d(x, \gamma)$, and so $x \in \Psi_\gamma$.

Recall now that, by Lemma 4, K_u is a ray or line, and note that $K_u \cap VR(\gamma, S)$ is not empty: it contains at least u and a small segment around it. Consider now two points $p, q \in K_u \cap VR(\gamma, S)$, and a point $y \in K_u$ between p and q . Let v be a point on $\bigcup\{s \in S \mid s \neq \gamma\}$ closest to y (v exists since $\bigcup\{s \in S \mid s \neq \gamma\}$ is compact). Since $p, q \in VR(\gamma, S)$, we have $d(p, u) \leq d(p, v)$ and $d(q, u) \leq d(q, v)$. Since $\{x \in \mathbb{R}^2 \mid d(x, u) \leq d(x, v)\}$ is a closed half-plane, this implies $d(y, u) \leq d(y, v) = \min_{s \in S \setminus \{\gamma\}} d(y, s)$, and so $y \in VR(\gamma, S)$. If y belongs to the closure of $VR(s, S)$ for some $s \in S \setminus \{\gamma\}$, then we have $d(y, v) = d(y, s) = d(y, u)$. So y belongs to the bisector of u and v , which implies that $d(p, u) \geq d(p, v)$ or $d(q, u) \geq d(q, v)$. Remember that $d(p, u) \leq d(p, v)$ and $d(q, u) \leq d(q, v)$, so p or q belongs to the bisector of u and v . Then the bisector of u and v is the line pq , and so $v = u$, a contradiction. It follows that $K_u \cap VR(\gamma, S)$ is a segment, and its interior points do not lie in the closure of any other Voronoi region.

Suppose now that h_ℓ is not continuous. Then there exists a sequence of points in γ converging to $u \in \gamma$ whose h_ℓ -values do not converge to $h_\ell(u)$. Since $VR(\gamma, S)$ is

bounded, there exists a subsequence converging to some $w \neq h_\ell(u)$. Because of the continuity of ψ_γ , $\psi_\gamma(w) = u$ and so, by Lemma 4, w lies on the straight line through u and $h_\ell(u)$. Since $\partial VR(\gamma, S)$ is closed, $w \in \partial VR(\gamma, S)$. This implies that w belongs to the closure of two different Voronoi cells, although it lies in the interior of $K_u \cap VR(\gamma, S)$, which is a contradiction to the above. \square

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

Theorem 8. *Let S be a harmless site collection.*

- (i) *The union of the Voronoi regions covers the interior of 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 from s .*
- (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.*
- (vi) *Any triple of sites defines at most two Voronoi vertices.*

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 $s' \in S$. Clearly this implies $x \in VR(s, R)$.

(iii) Let $x \in VR(s, S) \cap VR(t, S)$. There must be points $u \in s$, $v \in t$ lying on the maximal empty circle around x . By Lemma 7 any point in the interior of the line segments \overline{xu} and \overline{xv} 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) First assume that s is a curve. Let α be some closed curve within $VR(s, S)$. Let q be some arbitrary point surrounded by α and let $K := \{x \in VR(s, S) \mid x \equiv q\}$. K intersects α in at least two points $a, b \in VR(s, S)$ surrounding q , and so $q \in VR(s, S)$ by Lemma 7. 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.

Now assume that s is a point and $p \in VR(s, S)$. Let p' be a point in the interior of the segment ps . We denote by C (resp. C') the circle centered at p (resp. p') through s . Since the interior of C does not intersect any site, the same is true for C' , hence $p' \in VR(s, S)$.

(v) Let s be a curve, and consider the functions $h_\ell, h_r: s \mapsto \partial VR(s, S)$ defined in Lemma 7. For all $x \in s$, we have $h_\ell(x) \equiv h_r(x) \equiv x$. By Lemmas 6 and 7 we know that h_ℓ and h_r are homeomorphisms. Thus the images $A := h_\ell(s)$ and $B := h_r(s)$ are simple arcs. For 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 Fig. 7). Since I, A, J, B are pairwise nonintersecting, their concatenation forms a Jordan




Fig. 7. Voronoi region $VR(s, S)$.

curve. If s is a circular segment then one of A, B may degenerate to a single point, the center of s . 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 (Fig. 8(a), $VR(s, S)$ is the segment uv) or just s itself (Fig. 8(b)). 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.

(vi) Assume three sites $s_1, s_2, s_3 \in S$ define three Voronoi vertices v_1, v_2, v_3 . If some $VR(s_i, S)$ is a point or line segment then, by our definition, its boundary consists of at most two vertices, contradiction. Otherwise, it follows from (v) that each $VR(s_i, S)$ is bounded by a Jordan curve. For each pair (s_i, v_j) we can draw a path connecting s_i to v_j in the interior of $VR(s_i, S)$, so that these paths are disjoint. However, that is a planar embedding of $K_{3,3}$, a contradiction. \square

Theorem 8(iv) implies that the edges of the Voronoi graph are indeed well-defined and are curves. Finally, we consider the complexity of this Voronoi graph.

Theorem 9. *Given a harmless site collection S of n sites (not counting the outer circle ω). 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.*

Proof. Each face of the graph is a Voronoi region, and there are n of those plus the outer face. All Voronoi regions are simply connected, so $V(S)$ is connected. Let v and e denote the number of vertices and edges of the graph. By definition, a Voronoi vertex has degree at least three, so $2e \geq 3v$. By Euler's formula, we have $2 \leq n + 1 - e + v \leq n + 1 - e + 2e/3$, and so $e \leq 3n - 3$. That implies $v \leq 2n - 2$. \square




Fig. 8. Degenerate Voronoi-cells.

5. Randomized Incremental Construction

The Voronoi diagram of curved objects does not fit into the framework of abstract Voronoi diagrams by Klein et al. [15] 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: the bisector can be a closed curve as in Fig. 1.

If, however, the objects form a harmless site collection, 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 [6], Mulmuley [17], and Boissonnat et al. [3]. We have to make sure that during the execution of 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 the Voronoi diagram of n points $P \cup Q$. 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 in the relative interior of s .

The points in Q serve as “placeholders” for the curves they stem from. In the second stage the curve sites are added one by one in random order. The 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 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 [16], [20]. 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 3, 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, \dots, s_m be a random permutation of the curve sites in S . Let q_r denote 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 enlarge a point site q_r in the current Voronoi diagram $V(P \cup \{\omega, s_1, \dots, s_{r-1}, q_r, \dots, q_m\})$ to the curve site s_r and obtain $V(P \cup \{\omega, s_1, \dots, s_r, q_{r+1}, \dots, q_m\})$. Let us look at this in more detail.

For brevity, let $s := s_r$, $q := q_r$, $R := P \cup \{s_1, \dots, s_{r-1}, q_r, q_{r+1}, \dots, q_m\}$, and $R' := P \cup \{s_1, \dots, s_{r-1}, s_r, q_{r+1}, \dots, q_m\}$. By Theorem 8, 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 $VR(s, R')$. (See Fig. 9.) We first prove a lemma.

Lemma 10. *The “skeleton” \mathcal{I} (the portion of $V(R)$ that lies in $VR(s, R')$) 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.*




Fig. 9. The Voronoi diagram before and after inserting s . The shaded region in (a) is $VR(q, R)$ and the shaded region in (b) is $VR(s, R')$. The skeleton \mathcal{I} is represented by the dashed edges.

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 8(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 to Theorem 8(iv). 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. \square

Consider now the boundary of $VR(s, R')$. As discussed in the proof of Theorem 8(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 determine 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 10 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 [4], [21]: Fix R' , and let s be a random curve site in R' . The total complexity of $V(R')$ is $O(n)$ by Theorem 9, 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 $O(\sum_{r=1}^n n/r)$ which is $O(n \log n)$ as well.

Theorem 11. *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 thank Rida Farouki, Stefan Felsner, and Lorenz Wernisch for fruitful discussions and valuable hints. We also thank an anonymous referee for his or her helpful comments.

References

1. H. Alt and O. Schwarzkopf. The Voronoi diagram of curved objects. In *Proc. 11th Annual ACM Symposium on Computational Geometry*, pages 89–97, 1995.
2. F. Aurenhammer. Voronoi diagrams: a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, September 1991.
3. 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.
4. L. P. Chew. Building Voronoi diagrams for convex polygons in linear expected time. Technical Report PCS-TR90-147, Dept. Mathematics and Computer Science, Dartmouth College, Hanover, NH, 1986.
5. J. J. Chou. Voronoi diagrams for planar shapes. *IEEE Comput. Graph. Appl.*, 15:52–59, March 1995.
6. K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
7. G. Elber and M. Kim. Bisector curves of planar rational curves. *CAD*, 30:1089–1096, 1998.
8. R. T. Farouki and R. Ramamurthy. Degenerate point/curve and curve/curve bisectors arising in medial axis computations for planar domains with curved boundaries. *Comput. Aided Geom. Design*, 15:615–635, 1998.
9. R. T. Farouki and R. Ramamurthy. Voronoi diagram and medial axis algorithm for planar domains with curved boundaries, I. *J. Comput. Appl. Math.*, 102:119–131, 1999.
10. R. T. Farouki and R. Ramamurthy. Voronoi diagram and medial axis algorithm for planar domains with curved boundaries, II. *J. Comput. Appl. Math.*, 102:253–277, 1999.
11. H. W. Guggenheimer. *Differential Geometry*. McGraw-Hill, New York, 1963.
12. M. I. Karavelas and M. Yvinec. The Voronoi diagram of planar convex objects. In *Proc. 11th European Symposium on Algorithms (ESA 2003)*, pp. 337–348. Lecture Notes in Computer Science, vol. 2832. Springer-Verlag, Berlin, 2003.
13. R. Klein. Abstract Voronoi diagrams and their applications. In *Computational Geometry and its Applications*, pp. 148–157. Lecture Notes in Computer Science, vol. 333. Springer-Verlag, Berlin, 1988.
14. R. Klein. *Concrete and Abstract Voronoi Diagrams*. Lecture Notes in Computer Science, vol. 400. Springer-Verlag, Berlin, 1989.

15. R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom. Theory Appl.*, 3:157–184, 1993.
16. D. E. Muller and F. P. Preparata. Finding the intersection of two convex polyhedra. *Theoret. Comput. Sci.*, 7:217–236, 1978.
17. K. Mulmuley. A fast planar partition algorithm, I. In *Proc. 29th Annual IEEE Symposium on the Foundations of Computer Science*, pages 580–589, 1988.
18. C. Ó'Dúnlaing and C. K. Yap. A “retraction” method for planning the motion of a disk. *J. Algorithms*, 6:104–111, 1985.
19. A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley, Chichester, 1992.
20. F. P. Preparata and M. I. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.
21. R. Seidel. Backwards analysis of randomized geometric algorithms. In J. Pach, editor, *New Trends in Discrete and Computational Geometry*, pages 37–68. Algorithms and Combinatorics, vol. 10. Springer-Verlag, Berlin, 1993.
22. J. J. Stoker. *Differential Geometry*. Wiley-Interscience, New York, 1969.
23. W. Vogt. Über monotongekrümmte Kurven. *J. Reine Angew. Math.*, 144:239–248, 1914.
24. 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.
25. C. K. Yap and H. Alt. Motion planning in the CL-environment. In *Proc. 1st Workshop on Algorithms and Data Structures*, pp. 373–380. Lecture Notes in Computer Science, vol. 382. Springer-Verlag, Berlin, 1989.

Received July 20, 2004, and in revised form May 11, 2005. Online publication August 4, 2005.