Abstract Voronoi Diagrams Revisited

Rolf Klein* Elmar Langetepe* Zahra Nilforoushan†

Abstract

Abstract Voronoi diagrams [21] were designed as a unifying concept that should include as many concrete types of diagrams as possible. To ensure that abstract Voronoi diagrams, built from given sets of bisecting curves, are finite graphs, it was required that any two bisecting curves intersect only finitely often; this axiom was a cornerstone of the theory. In [12], Corbalan et al. gave an example of a smooth convex distance function whose bisectors have infinitely many intersections, so that it was not covered by the existing AVD theory. In this paper we give a new axiomatic foundation of abstract Voronoi diagrams that works without the finite intersection property.

Keywords: Abstract Voronoi diagrams, computational geometry, distance problems, Voronoi diagrams.

1 Introduction

Voronoi diagrams belong to the most interesting and useful structures in geometry. Dating back to Descartes [13], and known to mathematicians ever since (see, e. g., Gruber [16]), Voronoi Diagrams were the topic of a seminal paper by Shamos and Hoey [32] that helped creating a new field, computational geometry. The general idea is quite natural. There is a space on whom some objects, called sites, exert a certain influence. Each point of the space belongs to the region of that site whose influence is strongest. Most often influence is reciprocal to distance. Meanwhile, CiteSeer lists more than 4800 related articles on Voronoi diagrams. Surveys focussing on their structural and algorithmic aspects were presented by Aurenhammer [6], Aurenhammer and Klein [7], Fortune [15], and, for generalized Voronoi diagrams, by Boissonnat et al. [9]. Beyond their value to computer science, Voronoi diagrams have important applications in many other sciences; prominent examples can be found in Held [17] and in Okabe et al. [30].

For many years, computational geometers have studied Voronoi diagrams in the plane that differed by the types of sites and distance measures used. Typically, algorithms were hand-tailored to fit a particular setting. This situation called for a unifying view. An elegant structural approach was by Edelsbrunner and Seidel [14] who suggested to define general Voronoi diagrams as lower envelopes of suitable "cones". Independently, *Abstract Voronoi Diagrams (AVDs)* were introduced by the first author in [20], as a unifying concept for both, structure theory and algorithmic computation.

The basic observation behind AVDs was that Voronoi diagrams are built from systems of bisecting curves that have certain combinatorial properties in common, whereas the nature of the sites and of the distance function are of secondary importance. A challenge was in finding a small set of simple axioms for bisecting curve systems. They should ensure that a Voronoi diagram formed from such a curve system has desirable structural properties (like being a finite plane graph of linear complexity), and that it can be efficiently computed. At the same time, this approach should be as general as possible.

*University of Bonn, Institute of Computer Science I, D-53117 Bonn, Germany.
†Amirkabir University of Technology, Faculty of Mathematics and Computer Science, Tehran, Iran.

In order to achieve the goals just mentioned, AVDs were defined in the following way. For any two elements p, q of a set S of indices, also referred to as sites, a curve $J(p, q)$ was given, that splits the plane into two unbounded open domains. One of these domains was labeled by p, the other by q; these labels were part of the definition of $J(p, q)$. The curve itself was added to one of the two domains according to some global order $\prec$ on S. The Voronoi region of p was defined as the intersection of all sets associated with p; detailed definitions will be be given in Section 5.

Now three properties were required of the given curves and of order $\prec$. Voronoi regions should be path-connected, and their union should cover the whole plane. Moreover, any two curves $J(p, q), J(r, s)$ should intersect only finitely often. These requirements are met, for example, by the Euclidean Voronoi diagrams of points or line segments, additive weights, power diagrams, and all convex distance functions whose circles are semi-algebraic.

It turned out that these axioms were strong enough to ensure that abstract Voronoi diagrams have many of the properties found in diagrams based on concrete sites and distance functions, and that they can be constructed efficiently.

The finite intersection assumption was instrumental in analyzing the structure of abstract Voronoi diagrams. It was applied twice. First, in proving the topological fact that in a neighborhood of any point v, the bisecting curves passing through v form a star; see the "piece of pie" Lemma 2.3.2 [21]. This fact allowed a local view on which combinatorial definitions could be based. Second, the finite intersection property was explicitly used to guarantee that abstract Voronoi diagrams are finite planar graphs; see Lemma 2.4.2 [21].

Three asymptotically optimal AVD algorithms have been developed, each for a certain subclass of AVDs. A deterministic $O(n \log n)$ divide & conquer algorithm [21], based on work by Shamos and Hoey [32] and by Chew and Drysdale [10], for situations where recursive partitions with cycle-free bisectors are guaranteed; a deterministic linear time algorithm [23] for situations resembling "general convex position", based on the technique by Aggarwal et al. [2], and an $O(n \log n)$ randomized incremental construction algorithm [24] for AVDs whose regions have path-connected interiors, based on work by Clarkson and Shor [11].

McAllister et al. [5], Ahn et al. [3], Karavelas and Yvinec [19], Abellanas et al. [1], Aichholzer et al. [4], and Bae and Chwa [8] presented new types of Voronoi diagrams that were under the umbrella of the AVD concept. The notion of abstract Voronoi diagrams has been generalized to furthest site diagrams by Mehlhorn et al. [28], to dimension 3 by Lê [25], and to a dynamic setting by Malinauskas [26]. A slightly simplified version of abstract Voronoi diagrams has been implemented in LEDA by Seel [33].

But Corbalan et al. [12] gave an example of a convex distance function whose bisectors have an infinite number of intersections; its unit circle is smooth, but not semi-algebraic. The existing AVD concept, with its definitions and proofs relying on the finite intersection property, did not cover this example.

The purpose of this paper is in proving that abstract Voronoi diagrams can be defined and constructed without the finite intersection assumption. In fact, the other two axioms—that Voronoi regions be path-connected and cover the plane—are just strong enough to imply what is needed. Although the proof of this fact requires new techniques quite different from those used in [21, 24], we think this effort is well-invested. First, the class of concrete Voronoi diagrams covered by the AVD concept grows; in particular, all convex distance functions are included now. Second, with one axiom less to check, applying AVDs becomes easier. Third, there is scientific value (and aesthetic pleasure) in minimizing axiomatic systems.

Some care is necessary in dealing with general curves that can intersect each other infinitely often. To keep the analysis simple, we require, in this paper, that not only the Voronoi regions, but also their interiors, are path-connected. This was also postulated for the randomized incremental algorithm in [24]. It helped to avoid the complications in [21] that were caused by the fact that Voronoi edges and vertices could form connections between several parts of one Voronoi region. Our requirement also allows us to abandon the order $\prec$, which was used to distribute the bisecting curves among the sites.

The rest of this paper is organized as follows. In Section 3 we state the new set of axioms, and derive some preliminary facts. The main part is Section 4, where we show that AVDs based on the new axioms are finite plane graphs, without resorting to the finite intersection assumption. This is accomplished in the following way. First, we prove that a bisecting curve $J(p, q)$ cannot more than twice alternate between the domains separated by some curve $J(p, r)$, without disconnecting a Voronoi region; see Lemma 6.[1] This allows us to analyze how $J(p, q)$ and $J(p, r)$ can behave in the neighborhood of an intersection point, without having a "piece of pie" lemma available. For sets S of size 3 we show, in Lemma 8, that each point w on the boundary of a Voronoi region is accessible from this region. That is, there exists an arc α with endpoint w such that α without w is fully contained in the Voronoi region.[2]

Using an elegant argument by Thomassen [34], accessibility implies that an abstract Voronoi diagram of three sites contains at most two points that belong to the closure of all Voronoi regions. From this one can directly conclude that AVDs of many sites are finite plane graphs; see Theorem 10. Now a piece of pie lemma can be shown at least for the Voronoi edges meeting at a Voronoi vertex, which is sufficient for our purposes.

In Subsection 4.3 we show that a curve system for index set S fulfills our axioms iff this holds for each subset S' of size three. This fact was observed in [22] for the old AVD model; the proof given in Subsection 4.3 is new and more general.

In Section 5 we address the construction of abstract Voronoi diagram based on the new axioms. With the finite intersection assumption and order $\prec$ removed, the class of curve systems to which randomized incremental construction can be applied, is now strictly larger than in [24]. Divide & conquer can be applied if acyclic partitions are possible, as in [21]; but curve systems causing the interior of Voronoi regions to be disconnected are no longer admissible. By this restriction, the divide & conquer algorithm becomes considerably simpler.

2 Acknowledgement

The authors would like to thank the anonymous referees for their careful reading of the first version of this paper, and for their valuable comments and suggestions.

3 The new AVD axioms

We are given a finite set S of indices (representing sites), and, for any two indices $p \neq q$ of S, a curve $J(p, q) = J(q, p)$ that splits the plane into two unbounded domains, labeled $D(p, q)$ and $D(q, p)$.[3] These labels are assigned to the two domains as part of the definition of $J(p, q)$; see Figure 1. We define, for each $p \in S$, the set

$$\mathrm{VR}(p, S) := \bigcap_{q \in S \setminus \{p\}} D(p, q), \tag{1}$$

[1] One should observe that both curves are associated with the same site, p. In the old AVD model [21], this observation was an easy consequence of the fact that AVDs are finite plane graphs.

[2] By the Jordan curve theorem and its inverse, Jordan curves are characterized by accessibility; see Theorem 4.

[3] Informally, $D(p, q)$ and $D(q, p)$ will sometimes be called the "half-planes" defined by $J(p, q)$.

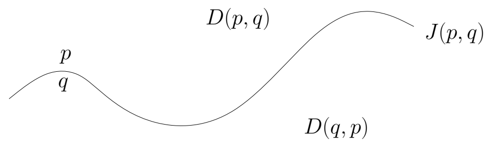

Figure 1: A bisecting curve.

and let

$$V(S) := \mathbb{R}^2 \setminus \bigcup_{p \in S} \mathrm{VR}(p, S). \tag{2}$$

Now we state which axioms the given curves must fulfill.

Definition 1 *The curve system* $J := \{J(p,q); p \neq q \in S\}$ *is called* admissible *if the following axioms are fulfilled.*

A_1) *Each curve* $J(p,q)$*, where* $p,q \in S$*, is mapped to a closed Jordan curve through the north pole by stereographic projection to the sphere.*[4]

A_2) *For each subset* $S' \subseteq S$ *and for each* $p' \in S'$*, the set* $VR(p', S')$ *is path-connected* [5]

A_3) *For each subset* $S' \subseteq S$*, we have*

$$\mathbb{R}^2 = \bigcup_{p' \in S'} \overline{VR(p', S')}.$$

Here, $\overline{A}$ denotes the topological closure of a set A in the Euclidean topology.

Definition 2 *For an admissible curve system* J *we call the set* $VR(p,S)$ *the* Voronoi region *of* p *with respect to* S*, whereas* $V(S)$ *is called the* Voronoi diagram *of* S*.*

Example. Figure 2 (i) shows an admissible curve system for $S = \{p, q, r\}$, and the resulting Voronoi diagram (ii). An index p placed closely to a bisecting curve $J(p,q)$ indicates on which side of $J(p,q)$ domain $D(p,q)$ is located.

In this example we observe some phenomena that cannot occur for Euclidean bisectors of points. There are points like a in the intersection of two bisecting curves $J(p,r)$ and $J(q,r)$ that do not lie on the third curve, $J(p,q)$. A point like w that is included in all three bisecting curves need not be a Voronoi vertex. The intersection of two bisecting curves, like $J(p,r)$ and $J(q,r)$, consists of an infinite number of connected components, curve segments or single points. These components may have accumulation points. In fact, in Figure 2 there is an infinite sequence of intersection points $a_i \in J(p,r) \cap J(q,r)$ that converge towards $a \in J(p,r) \cap J(q,r)$, such that each segment of $J(q,r)$ between a_i and a_{i+1} is disjoint from $J(p,r)$. Therefore, one must be careful not to speak of "the first point of $J(q,r)$ on $J(p,r)$ to the left of point a", etc..

[4]More precisely, the projected image is continuously completed, by the north pole, to a closed Jordan curve.

[5]We need not distinguish between path-connectedness and arc-connectedness because the Euclidean plane is Hausdorff. Thus, two points of a path-connected set can be connected not only by a path, which is a continuous image of $[0,1]$, but even by an *arc* which is image of $[0,1]$ under a homeomorphism, that is, of a bijective, bi-continuous mapping.

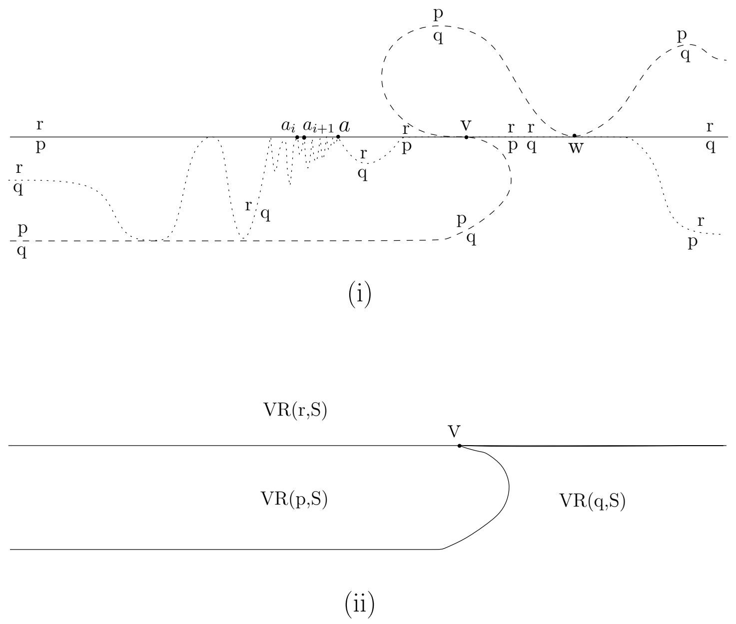

Figure 2: An admissible curve system (i) and the resulting Voronoi diagram (ii).

3.1 Preliminaries

In this section, and in the following one, we assume that J is an admissible curve system fulfilling axioms A_1, A_2, A_3 of Definition 1, unless stated otherwise.

First, we observe that the Voronoi diagram $V(S)$ can also be characterized in the following way; compare Lemma 2.2.1 in [21].

Lemma 3 *Let J be a system of admissible curves for index set S. Then,*

$$
\begin{aligned}
V(S) &= \bigcup_{p \neq q \in S} \overline{VR(p,S)} \cap J(p,q) \\
&= \bigcup_{p \neq q \in S} \overline{VR(p,S)} \cap \overline{VR(q,S)}
\end{aligned}
$$

Proof. "$\subseteq$:" If $z \in V(S)$ then $z \in \overline{\mathrm{VR}(p,S)} \setminus \mathrm{VR}(p,S)$ for some $p \in S$, by axiom A_3 and the definition of $V(S)$. Hence, there exists a site $q \neq p$ such that $z \notin D(p,q)$. Assume $z \in D(q,p)$. As this set is open, it would contain a whole neighborhood of z. But this contradicts $z \in \overline{\mathrm{VR}(p,S)} \subset \overline{D(p,q)}$. Therefore, $z \in J(p,q)$, thus proving the upper inclusion. The lower one is shown by contradiction. If no set $\overline{\mathrm{VR}(q,S)}$, where $q \neq p$, contained z then, by finiteness of S, for a neighborhood $U(z)$,

$$
U(z) \subseteq \bigcap_{q \neq p} \overline{\mathrm{VR}(q,S)}^{\mathrm{c}} = \left(\bigcup_{q \neq p} \overline{\mathrm{VR}(q,S)} \right)^{\mathrm{c}} \subseteq \overline{\mathrm{VR}(p,S)}
$$

would hold, contradicting $z \notin \mathrm{VR}(p,S)$.
"$\supseteq$:" Because of

$$
\overline{\mathrm{VR}(p,S)} \cap \overline{\mathrm{VR}(q,S)} \subseteq \overline{D(p,q)} \cap \overline{D(q,p)} \subset J(p,q)
$$

we need to consider only the upper inclusion. Let us assume that

$$z \in \overline{\mathrm{VR}(p,S)} \cap J(p,q).$$

If z were contained in some Voronoi region $\mathrm{VR}(r,S)$ then it would be an interior point of this region. Because of $z \in \overline{\mathrm{VR}(p,S)}$ this would imply $r = p$. But if z lies in the interior of the Voronoi region $\mathrm{VR}(p,S)$, it cannot be situated on $J(p,q)$, contradicting our assumption. Hence, $z \in V(S)$. □

In order to prove the finiteness of the Voronoi diagram as a graph, we shall employ some properties of plane curves that are sometimes stated as part of the Jordan curve theorem.

Theorem 4 *Let C be a plane curve, homeomorphic to a circle. Then $\mathbb{R}^2 \setminus C$ consists of two domains D_1, D_2 with common boundary C. For each point $z \in C$ there exists a neighborhood U whose boundary is homeomorphic to a circle, such that $U \setminus C$ consists of exactly two connected components. Point z is* accessible *from each domain $U \cap D_i$, that is, for each $p \in U \cap D_i$ there exists an arc α from p to z such that α minus its endpoint z belongs to $U \cap D_i$.*

Theorem 4 is a direct consequence of the stronger Jordan-Schönflies theorem, which states that a homeomorphism between a circle and a closed curve, C, in the plane can be extended to the whole plane, such that the interior of the circle is mapped onto the interior domain of C, and the circle's exterior to the exterior of C; cf. Rinow [31], for example. Theorem 4 also holds for closed curves on the sphere. Hence, it holds for the bisecting curves $J(p,q)$ we are dealing with, because they are mapped to Jordan curves through the north pole under stereographic projection. For simplicity, a homeomorphic image of the circle, or a homeomorphic image of the line that bisects the plane, will be called a *Jordan curve* in the sequel. As a trivial consequence of Theorem 4, *every* neighborhood of a point z on a Jordan curve C contains points of both domains D_i.

It is interesting to observe that the converse of Theorem 4 is also true. If C is a compact set whose complement in the plane consists of two connected components, such that each point of C is accessible from both, then C is a closed Jordan curve; see Thomassen [34] for a simple proof.

The following transitivity lemma will be a handy tool. Its claim would be trivial if we could read $z \in D(p,q)$ as "z is closer to p than to q". For the old AVD model [21], a similar statement with a slightly different proof was made in Lemma 3.5.1.1.

Lemma 5 *Let $p, q, r \in S$. Then $D(p,q) \cap D(q,r) \subseteq D(p,r)$ holds.*

Proof. Let $z \in D(p,q) \cap D(q,r)$. Point z must be contained in one of $D(r,p), J(r,p), D(p,r)$. If z were contained in $D(r,p)$, it could not lie in *any* of the closed Voronoi regions

$$\begin{aligned}
\overline{\mathrm{VR}(q,S')} &\subseteq \overline{D(q,p)} = D(q,p) \cup J(q,p) \\
\overline{\mathrm{VR}(r,S')} &\subseteq \overline{D(r,q)} = D(r,q) \cup J(r,q) \\
\overline{\mathrm{VR}(p,S')} &\subseteq \overline{D(p,r)} = D(p,r) \cup J(p,r)
\end{aligned}$$

for $S' := \{p,q,r\}$. This is impossible since these three sets cover $\mathbb{R}^2$ by axiom A_3.

Suppose $z \in J(r,p)$. By Theorem 4, there exists an arc α with endpoint z such that $\alpha \setminus \{z\} \subset D(r,p)$. With z, even a neighborhood U of z is contained in the open set $D(p,q) \cap D(q,r)$. Inside U, path α contains a point $z' \in D(p,q) \cap D(q,r) \cap D(r,p)$, which leads to the same contradiction as before. Consequently, the third case applies, that is, $z \in D(p,r)$. □

We note that Lemma 5 neither holds for the closures of the sets $D(\ ,\ )$, nor for the bisecting curves themselves.

4 The graph structure of $V(S)$

The main goal of this section is in proving that $V(S)$, where $|S| = n$, is a finite, plane graph with $\leq n$ faces, even though our bisecting curves do not fulfill the finite intersection property. To this end, we consider first an abstract Voronoi diagram of three sites, and show, in Lemma 8, that each point w on a region boundary is accessible from this Voronoi region. As we do not have a piece of pie lemma available, which would grant us a clear view to a neighborhood of w, this requires some local analysis on the bisecting curves passing through w. This analysis will be based on the following Lemma 6, supported by Lemma 7.

Then the proof proceeds as follows. By Lemma 3, each point of $V(S)$ lies on the boundaries of at least two regions. We show that only finitely many points can be situated on three or more region boundaries; see Lemma 9. From this fact, the finiteness of $V(S)$ will be derived in Theorem 10.

4.1 Three sites

Let us consider the Voronoi region of a site p in the diagram $V(S)$ where $S = \{p, q, r\}$ consists of only three sites. For convenience we may assume that $J(p, q)$ is a horizontal line, and that $D(p, q)$ equals the lower half plane. The following lemma states that $J(p, r)$ can change at most twice between $D(p, q)$ and $D(q, p)$.

Lemma 6 *There cannot be four points consecutively visited by $J(p, r)$ that belong alternately to $D(p, q)$ and $D(q, p)$.*

Proof. Suppose that our claim is wrong, and that $J(p, r)$ does visit four points a_1, a_2, a_3, a_4, in this order, such that $a_1, a_3 \in D(q, p)$ and $a_2, a_4 \in D(p, q)$. The following facts will be helpful in deriving a contradiction. By π° we denote the relative interior of an arc π, that is, the arc without its endpoints.

Facts. We can connect

1. points a_2 and a_4 by an arc π such that π° is contained in $\mathrm{VR}(p, S)$,

2. points a_2 and a_4 by an arc ρ such that $\rho^\circ \subset \mathrm{VR}(r, S)$, and

3. points a_1 and a_3 by an arc σ such that $\sigma^\circ \subset \mathrm{VR}(q, S)$,

see Figure 4.

Proof. 1.) By Theorem 4, point a_2 is accessible by an arc α_2 from $D(p, r)$. Since a_2 is an interior point of $D(p, q)$, α_2 can be shortened to lie in $D(p, q)$. Then, α_2, without its endpoint a_2, is contained in $D(p, q) \cap D(p, r) = \mathrm{VR}(p, S)$. Similarly, there is an arc α_4 accessing point a_4 from $\mathrm{VR}(p, S)$. W. l. o. g., α_2 and α_4 are disjoint. Their respective endpoints a_2' and a_4' in $\mathrm{VR}(p, S)$ can be connected by an arc α entirely running in $\mathrm{VR}(p, S)$, by axiom A_2. Should α_2 intersect α in a point different from a_2', let a_2'' denote the first point of α met when traversing α_2 from a_2 towards a_2'.[6] We cut both α_2 and α at point a_2'', and perform similar surgery on α_4, if necessary. The concatenation of the three resulting arcs yields an arc with the properties desired; see Figure 3[7].

[6] This point is well-defined. Indeed, if $f(t)$ is a parametrization of α_2 satisfying $f(0) = a_2$, then $t'' := \sup\{t \geq 0; f([0, t]) \subset \alpha^c\}$ exists, and $f(t'') = a_2''$ holds.

[7] One should observe that in Figures 3 and 4 domain $D(r, p)$ is depicted to be above curve $J(p, r)$. Our proof does not make use of such an assumption.

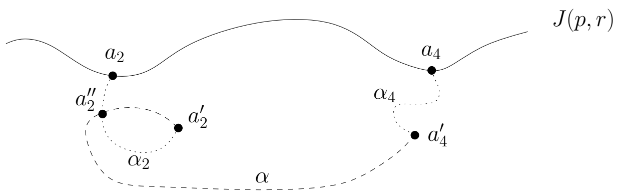

Figure 3: Constructing an arc that runs through $\mathrm{VR}(p, S)$ and connects a_2 to a_4.

2.) Point a_2 is also accessible from $D(r, p)$ by an arc β_2. Clearly, β_2° lies in $D(r, p) \cap D(p, q) \subseteq D(r, q)$, by Lemma 5, hence in $\mathrm{VR}(r, S)$. The rest of the proof is analogous to the proof of (1).

3.) Point a_1 is accessible from $D(p, r)$ by an arc γ_1 contained in $D(q, p)$. Thus, γ_1° is contained in $D(q, p) \cap D(p, r) \subseteq D(q, r)$, hence in $\mathrm{VR}(q, S)$. The same holds for a_3, and we continue as before. This concludes the proof of the three facts. $\square$

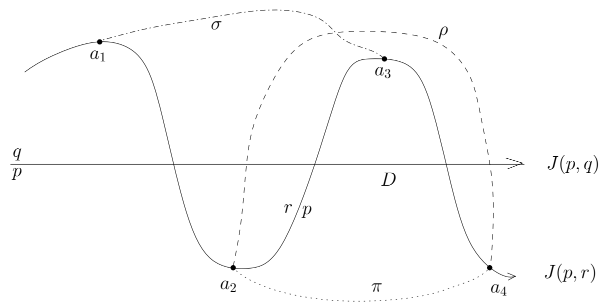

Figure 4: Domain D contains point a_3, but not a_1. Hence, path σ must intersect either ρ or π, in order to connect a_1 and a_3. Both alternatives are impossible because all arcs are contained in different Voronoi regions.

To complete the proof of Lemma 6, we argue as follows. Together, arcs $\pi \subset \mathrm{VR}(p, S)$ and $\rho \subset \mathrm{VR}(r, S)$ form a closed Jordan curve; let D denote its interior domain, as shown in Figure 4.

We observe that $J(p, r)$ cannot intersect the relative interiors of either path π or ρ, because these are contained in $D(p, r)$ and $D(r, p)$, respectively. Being a simple curve, $J(p, r)$ can pass through points a_2 and a_4 only once. On the other hand, $J(p, r)$ must pass through D to separate π from ρ. Therefore, the segment of $J(p, r)$ between a_2 and a_4 is fully contained in domain D, while the two unbounded complementary segments of $J(p, r)$ stay outside. Consequently, we have $a_3 \in D$ and $a_1 \in D^c$. But the path σ connecting a_1 to a_3 is contained in $\mathrm{VR}(q, S)$ and, therefore, unable to intersect the boundary of D, which belongs to the closures of the regions of p and r. This contradiction completes the proof of Lemma 6. $\square$

The next step is in proving that, for Voronoi diagrams of three sites, each point w on the boundary of a Voronoi region is accessible from this region; see Lemma 8 below. To this end, we need to discuss the different ways in which two bisecting curves $J(p, q)$ and $J(p, r)$ can intersect at some point w. Let $g(t)$ be a parametrization of $J(p, r)$ such that $D(p, r)$ lies on the right hand side of $J(p, r)$ as t tends to $+\infty$. Suppose that $g(0) = w$ holds. As t approaches 0 from below, the points $g(t)$ cannot alternate between $D(p, q)$ and $D(q, p)$ infinitely often, thanks to Lemma 6. Thus, there exist $\delta, \delta' > 0$ such that $G^- := g((-\delta, 0))$ is included in $\overline{D(p, q)}$ or in $\overline{D(q, p)}$. The same holds for $G^+ := g((0, \delta'))$. Analogously, there are two segments F^-, F^+ of $J(p, q)$ before and after w, each of which is contained in one of the sets $\overline{D(p, r)}$ or $\overline{D(r, p)}$; see

Figure 5 for an example. Let us first assume that none of the segments G^-, G^+ is fully contained

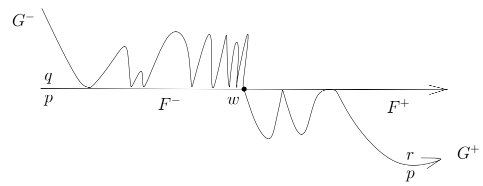

Figure 5: Here G^- stays in $\overline{D(q,p)}$, while G^+ runs in $\overline{D(p,q)}$. Moreover, $F^- \in \overline{D(p,r)}$ and $F^+ \in \overline{D(r,p)}$.

in $J(p,q)$, and none of F^-, F^+ in $J(p,r)$.[8] To facilitate our analysis, we first consider curve $J(p,q)$ and distinguish between the following cases.

$$\begin{array}{lcl} A & : & F^-,\ F^+ \subset \overline{D(p,r)} \\ B & : & F^- \subset \overline{D(p,r)} \text{ and } F^+ \subset \overline{D(r,p)} \\ C & : & F^- \subset \overline{D(r,p)} \text{ and } F^+ \subset \overline{D(p,r)} \\ D & : & F^-,\ F^+ \subset \overline{D(r,p)} \end{array}$$

Analogous cases A, B, C, D are possible for $J(p,r)$; they result from replacing r with q, and F with G, in this definition. In principle, 16 combinations AA, AB, AC,..., DC, DD of these cases should be considered. However, we observe that XY corresponds to YX under the symmetry $\Sigma : r \leftrightarrow q, F \leftrightarrow G$. Of the remaining 10 combinations, only 4 are geometrically possible, as we shall conclude from the following lemma. Intuitively, it states that the facts "J touches J' at w" and "J crosses J' at w" are symmetric in J and J'.

Lemma 7 *(i) If G^-, G^+ are contained in the closure of the same half-plane defined by $J(p,q)$, the same holds for F^-, F^+ with respect to $J(p,r)$.*
(ii) If $F^- \subset \overline{D(p,r)}$ and $F^+ \subset \overline{D(r,p)}$ then $G^- \subset \overline{D(q,p)}$ and $G^+ \subset \overline{D(p,q)}$. The same holds with $+$ and $-$ reversed.

Proof. (i) Let U be a neighborhood of w, chosen by Theorem 4, such that $J(p,r) \cap U$ is one connected segment contained in the union of G^-, G^+, and $\{w\}$. We can make U small enough to guarantee that the part of $J(p,q)$ passing through U is also contained in F^-, F^+, and $\{w\}$, but possibly disconnected; see Figure 6. By way of contradiction, assume that G^-, $G^+ \subset \overline{D(p,q)}$, but that there are points $z \in D(p,r) \cap F^-$ and $z' \in D(r,p) \cap F^+$ close to w on $J(p,q)$. The segment H of $J(p,q) \cap U$ that contains z, w, z', divides domain U into two domains, U_1 to the left of H, and U_2 to the right of H. Both have Jordan curves as boundaries. Since z, z' are accessible from U_1, there exists an arc $\alpha_1 \subset U_1$ connecting them. We can assume that α_1 stays in the open half-plane $D(q,p)$. Namely, each of the (at most countably many) excursions of α_1 to $D(p,q)$ can be replaced with circular arcs in $D(q,p)$, as depicted in Figure 6.[9] Since arc α_1 connects points z, z' from both sides of $J(p,r)$, it must meet $J(p,r)$ at some point $y \in U_1 \cap D(q,p)$. But each point $y \neq w$ of $J(p,r)$ in U belongs to G^- or to G^+, which are contained in the closure of $D(p,q)$, by assumption. Contradiction!

[8] Under this assumption, $G^- \subset \overline{D(p,q)}$ implies the following. Moving along G^- towards w, one never enters $D(q,p)$. One always meets another point of $D(p,q)$, and perhaps points of $J(p,q)$ in between. The latter may accumulate.

[9] Only finitely many circular arcs are needed to this end. This can be seen as follows. Let I denote the interval of $J(p,q)$ connecting an exit and re-enter point of α_1. For each point of I a circular neighborhood is contained in U_1. Since I is compact, finitely many of these circles cover I.

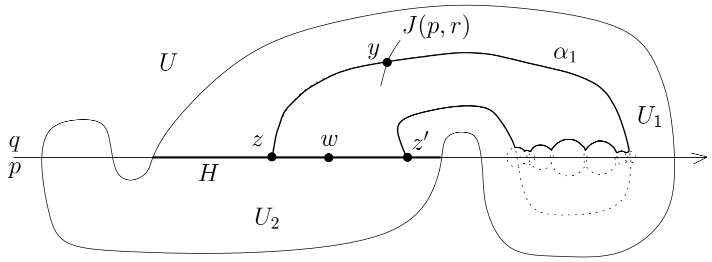

Figure 6: Illustrating the proof of Lemma 7.

(ii) Now assume $F^- \subset \overline{D(p,r)}$ and $F^+ \subset \overline{D(r,p)}$, and let $z \in D(p,r) \cap F^-$ and $z' \in D(r,p) \cap F^+$ as in the proof of (i). In addition to α_1, there exists an arc $\alpha_2 \subset U_2 \cap D(p,q)$ connecting z to z'. Curve $J(p,r)$ must intersect both arcs, α_1 and α_2; hence, it must visit both $D(q,p)$ and $D(p,q)$. Since both G^- and G^+ are fully contained in the closure of one of these half-planes, there are but two possibilities. Either $G^- \subset \overline{D(q,p)}$ and $G^+ \subset \overline{D(p,q)}$, which is what we claim, or $G^- \subset \overline{D(p,q)}$ and $G^+ \subset \overline{D(q,p)}$, which leads to a contradiction, because $z \in D(p,r)$ cannot be situated to the left of $J(p,r)$. □

Clearly, Lemma 7 remains true under symmetry Σ introduced before Lemma 7. A quick inspection shows that this leaves us with only 4 combinations, namely AD, DD, AA, and BC. They are illustrated in Figure 7.

It remains to account for those cases where the curves $J(p,q), J(p,r)$ share one or two segments close to their intersection point w. We shall now demonstrate how to view all situations possible as special cases of the configurations AD, DD, and BC displayed in Figure 7. The equality and unequality signs shown in this figure indicate which curve segments may coincide and which are supposed to be different. The case analysis given in Lemma 8 will be in accordance with these properties.

The situations where both G^- and G^+ are contained in $\overline{D(q,p)}$ are included in AD or DD, respectively, depending on the orientation of G. Otherwise, one of the segments G^-, G^+ lies in $\overline{D(p,q)}$ but not in F, while the other segment is part of F. We consider both cases in turn.

If $G^- \subset \overline{D(p,q)}$ and $G^+ \subset F$, two subcases are possible: $G^+ = F^-$, which reduces to the subcase ($F^+ = G^-$ and $F^- \subset \overline{D(p,r)}$) of BC under symmetry Σ, and $G^+ = F^+$, which reduces, under Σ, to the subcase ($F^+ = G^+$ and $F^- \subset \overline{D(p,r)}$) of AD.

If $G^+ \subset \overline{D(p,q)}$ and $G^- \subset F$, we have to consider the subcase $G^- = F^-$, which reduces, under Σ, to the subcase ($F^- = G^-$ and $F^+ \subset \overline{D(p,r)}$) of AD, and the situation where $G^- = F^+$, which is itself a subcase of BC.

Now we are ready to prove the main result of this subsection.

Lemma 8 *Let w be a point on the boundary of VR(p,S), where $|S| = 3$. Then there exists an arc α with endpoint w such that $\alpha \setminus \{w\} \subset$ VR(p,S) holds.*

Proof. First, we assume that w is contained in only one bisecting curve, $J(p,r)$. For w to belong to the boundary of the region of p, it cannot be in $D(q,p)$. Thus, $w \in D(p,q)$. By Theorem 4, applied to $J(p,r)$ on the sphere, there is an arc α accessing w from $D(p,r) \cap D(p,q) = \mathrm{VR}(p,S)$.

Now we assume that w is contained in both, $J(p,q)$ and $J(p,r)$. By the previous discussion, we need only inspect the four cases sketched in Figure 7.

AD) Here an arc accessing w from $D(p,q)$ is also contained in $D(p,r)$, after shortening, hence in $\mathrm{VR}(p,S)$.

DD) In this case $D(p,q) \cap D(p,r)$ is empty, in a neighborhood of w, so that point w cannot be on the boundary of $\mathrm{VR}(p,S)$, in contradiction to our assumption.

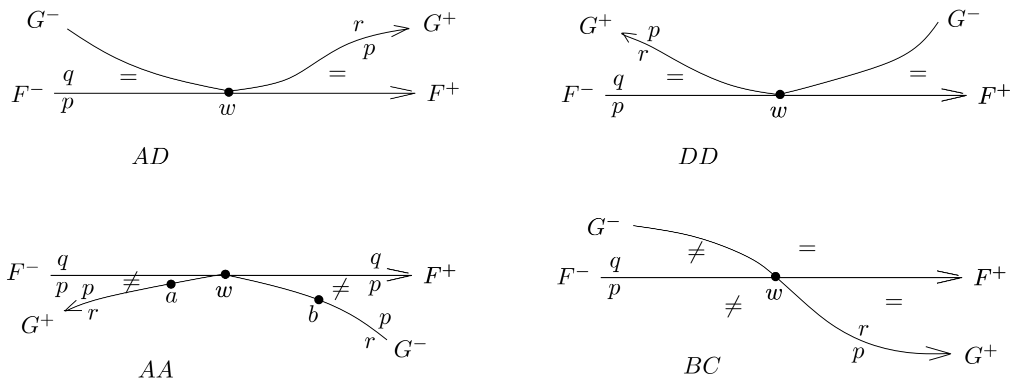

Figure 7: Four ways for $J(p,q), J(p,r)$ to pass through a point w.

AA) Here the Voronoi region of p would be disconnected. Formally, we could find points $a, b \in D(p,q)$ on G^+ and G^-, respectively, and connect them with paths $\pi \subset \mathrm{VR}(p,S)$ and $\rho \subset \mathrm{VR}(r,S)$ as in the proof of Lemma 6. The arguments presented there show that the domain D bounded by these paths must contain point w, so that ρ has to pass above w while π stays below $J(p,q)$. This is in conflict with the supposed orientation of $J(p,r)$ at w.
BC) This is the most interesting case. Let us walk along G^+ backwards, towards w. Suppose that we encounter two points b, d of F^- that appear in the order (d,b) on F^-, but in the order (b,d) on G^+; see Figure 8 (2). Then the Voronoi region of p would be disconnected, by the same formal proof as for AA. Thus, all points of F^- encountered on the backward walk along G^+ towards w, must be situated in the same order on both oriented segments, F^- and G^+, as shown in Figure 8 (1). This means, each new point of F^- we meet must be to the left of its predecessor on F^-, and, therefore, farther away from w. Consequently, a subsegment of G^+ with endpoint w must be wholly disjoint from F^-; let us denote it by G^+ again.

Now let U be a neighborhood of $w \in J(p,r)$ according to Theorem 4, small enough to intersect only the segments F^-, F^+ and G^-, G^+ of $J(p,q)$ and $J(p,r)$. Let H denote the segment of $U \cap F^-$ adjacent to w; see Figure 8 (3). Since $H \cap G^+ \subseteq F^- \cap G^+ = \emptyset$, segment H, together with parts of G^+ and ∂U, form a simple Jordan curve encircling a domain D, which is shaded grey in Figure 8 (3). Boundary point w is accessible via some arc α, where $\alpha^\circ \subset D$. Our assumption $F^- \subset \overline{D(p,r)}$ implies that H cannot enter $D(r,p)$; hence, D belongs to $U \cap D(p,r)$, implying $\alpha^\circ \subset D(p,r)$. On the other hand, at least a segment of α° starting from w must also belong to $D(p,q)$. This is because α can leave $D(p,q)$ only through points on $J(p,q)$ to the left of H. Therefore, part of α° lies in $\mathrm{VR}(p,S)$ and is an access path for w. □

4.2 Many sites

By Lemma 3, the Voronoi diagram $V(S)$ consists of all points in the plane that are contained in the closures of at least two Voronoi regions. Now we show that only finitely many points are contained in the boundary of three or more regions.

Lemma 9 *Let B be the set of all points on the boundaries of at least three Voronoi regions in some Voronoi diagram $V(S')$, where $S' \subseteq S$. If $|S| = n \geq 3$ then B is of size at most $2\binom{n}{3}$.*

Proof. First, we consider the case where $S = \{p_1, p_2, p_3\}$. If one of the three Voronoi regions is empty, $V(S)$ equals a single bisecting curve, and B is empty. Otherwise, we apply the following elegant argument of Thomassen's [34]. In each Voronoi region $\mathrm{VR}(p_j, S)$ we choose a point

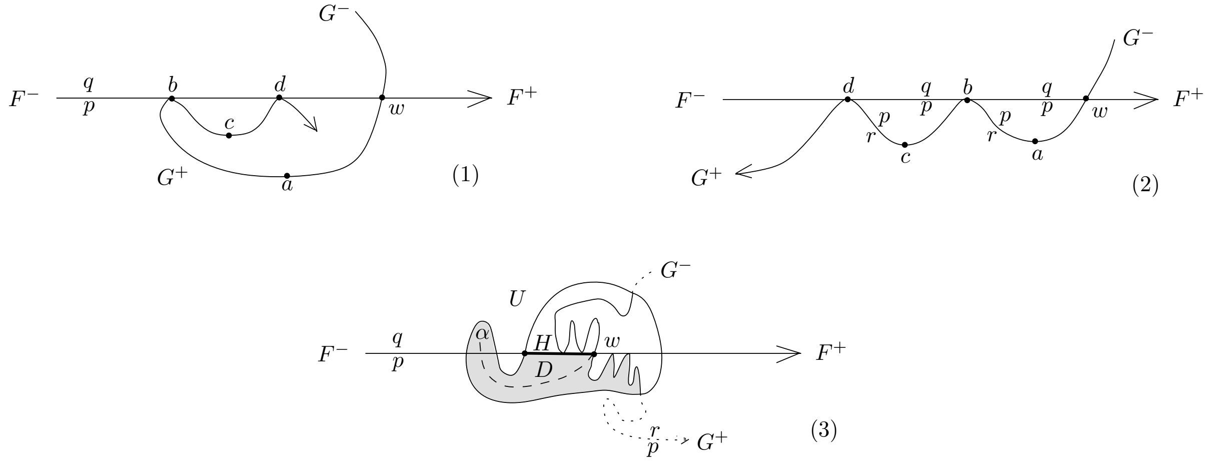

Figure 8: Discussion of case BC depicted in Figure 7.

denoted by a_j. Now suppose, by way of contradiction, that B contains three different points, v_1, v_2, v_3. By Lemma 8 and axiom A_2, we can find arcs $\alpha_{i,j}$ connecting v_i to a_j, such that $\alpha_{i,j} \setminus \{v_i\}$ is fully contained in $\mathrm{VR}(p_j, S)$. We may even assume that for each j, the three arcs $\alpha_{1,j}, \alpha_{2,j}, \alpha_{3,j}$ contained in the region of p_j form a plane tree T_j rooted at a_j.[10] Since the Voronoi regions of p_1, p_2, p_3 are disjoint, the trees T_1, T_2, T_3 are realizing a plane embedding of the bipartite graph $K_{3,3}$, which is impossible. Hence, $|B| \leq 2$ holds.

Now let $|S| > 3$, and let $v \in B$ be a point on the boundary of the Voronoi regions of $p, q, r \in S' \subseteq S$. Since the Voronoi regions of these indices in $V(S'')$, where $S'' := \{p, q, r\}$, can only be larger than those in $V(S')$, point v belongs to the region boundaries of p, q, r in $V(S'')$, too. We have just shown that there are at most two such points in $V(S'')$. Since there are only $\binom{n}{3}$ subsets S'' of S of size 3, the claim follows. □

We recall that a finite plane graph is an embedding in the plane of a finite abstract graph, that maps each edge e onto an arc whose endpoints are the embedded vertices adjacent to e. Two arcs do not intersect except at a common endpoint. If the two endpoints of an edge coincide, the edge is a loop, mapped onto a closed Jordan curve. Owing to the structure of Voronoi diagrams, we allow for a special vertex ∞, the inverse image of the north pole under stereographic projection to the sphere.[11]

Now we can prove the main result of this section.

Theorem 10 *The abstract Voronoi diagram $V(S)$, where $|S| = n$, is a finite plane graph of $O(n)$ edges and vertices.*

Proof. If $n = 2$ then $V(S)$ consists of a single bisecting curve with both endpoints at ∞, and we are done. Let us assume $n \geq 3$. By Lemma 3, the set $V(S)$ consists of all points contained in the closures of two or more Voronoi regions. The points on the boundary of three or more regions of $V(S)$ form a subset B' of B, which is finite by Lemma 9. The rest of $V(S)$ is decomposed into sets $B_{p,q}$ consisting of all points on the boundary of exactly the Voronoi regions $\mathrm{VR}(p, S)$

[10] This can be achieved as follows. First, we choose $\alpha_{1,j}$. Then we trace $\alpha_{2,j}$ from v_2 to a_j, and cut it at the first point of $\alpha_{1,j}$ it meets. Finally we trace $\alpha_{3,j}$ from v_3 to a_j, and cut it at the first point of the partial tree already constructed.

[11] Alternatively, one could compactify the Voronoi diagram in the way suggested in [21], by clipping the unbounded pieces of $V(S)$ at a sufficiently large circle.

and $\mathrm{VR}(q, S)$. Such a set $B_{p,q}$, if non-empty, consists of disjoint segments of the bisecting curve $J(p, q)$. We claim that each endpoint, v, of such a segment e belongs to the set $B' \cup \{\infty\}$.

This can be seen as follows. Suppose $v \neq \infty$. Clearly, v belongs to the closures of the regions of both, p and q, but the extension of e beyond v on $J(p, q)$ does not. Thus, we can either find points of some Voronoi region $\mathrm{VR}(r, S)$, where $r \notin \{p, q\}$, arbitrarily close to v on $J(p, q)$. Or a segment e' of $J(p, q)$ extending e beyond v belongs to $V(S)$. By Lemma 3, segment e' is contained in the closures of the Voronoi regions of two sites, at least one of which, r, is not in $\{p, q\}$. In either case, v in $\overline{\mathrm{VR}(r, S)}$.

Since $J(p, q)$ is simple, at most two segments e of $B_{p,q}$ can share a point of B'. Thus, the number of these segments is finite. We conclude that

$$V(S) \cup \{\infty\} = \bigcup_{p \neq q \in S} B_{p,q} \ \cup \ B' \ \cup \{\infty\}$$

is a finite plane graph. The elements of B' are the finite Voronoi vertices, while the edges are the segments of the sets $B_{p,q}$. Because $B_{p,q}$ and $B_{p',q'}$ are disjoint if $\{p, q\} \neq \{p', q'\}$, edges can intersect only at their endpoints. The $O(n)$ bound follows from the Euler formula, as usual. □

Once $V(S)$ is known to be a finite plane graph, the following "piece of pie" lemma can be shown; for a proof see, e. g., 40.20 in Rinow [31], or compare Lemma 2.3.1 in [21].

Lemma 11 *For each point v in the plane there exists an arbitrarily small neighborhood U of v, whose boundary is a simple closed curve, such that the following holds for each subset S' of S. Let $v \in V(S')$. If v is interior point of some Voronoi edge $e \subset B_{p,q}$ of $V(S')$ then U is divided by e in exactly two domains, one contained in $VR(p, S')$, the other in $VR(q, S')$. Otherwise v is a Voronoi vertex of $V(S')$, of degree $k \geq 3$. After suitably renumbering S', the Voronoi edges e_i adjacent to v belong to $B_{p_i,p_{i+1}}$ in counterclockwise order, where $0 \leq i \leq k-1$ is counted mod k. The edges e_{i-1} and e_i, together with ∂U, bound a piece of pie contained in $VR(p_i, S')$; these pieces are domains with Jordan curve boundaries. The sites $p_0, p_1, \ldots, p_{k-1}$ are pairwise different.*

The last fact can be seen as follows. Point v does not belong to any Voronoi region and can, therefore, not form a connection between different pieces of the same Voronoi region. On the other hand, no connecting path can "run around" some other piece, because abstract Voronoi regions are connected and their closures are simply-connected. This follows from Lemma 2.2.4 in [21]. For convenience, we include this statement and its simple proof.

Lemma 12 *Let $C \subset \overline{VR(p, S)}$ be a closed curve. Then each bounded connected component of $\mathbb{R}^2 \setminus C$ is contained in $VR(p, S)$.*

Proof. The complement of C consists of disjoint connected components exactly one of which, Z_∞, is unbounded. Let Z be a bounded connected component, and assume that some point $z \in Z$ does not belong to the Voronoi region of p. Then $z \in \overline{D(q, p)}$, for some $q \neq p$. We may even assume $z \in D(q, p)$ because a small enough neighborhood U of z is contained in the open set Z, and has nonempty intersection with $D(q, p)$, so that we could pick a suitable z' from U. Since $D(q, p)$ is unbounded, it contains a point $y \in Z_\infty$. Because $D(q, p)$ is path-connected there is an arc $\alpha \subset D(q, p)$ running from z to y. It must meet the curve C, which contradicts $C \subset \overline{D(p, q)}$. □

One should observe that from the axioms stated in Definition 1, only A1 was used in the proof of Lemma 12.

4.3 Characterizing admissible curve systems

In this subsection we show that only the subsets $S' \subseteq S$ of size 3 need to be checked, in order to ensure that a given curve system $J = \{J(p,q); p \neq q \in S\}$ is admissible in the sense of Definition 1. For the old AVD model, this fact has been stated in [22], and been proven for all bisecting curve systems whose (finitely many) pairwise intersections are proper crossings only. Here we give a different, more general proof based on Lemmata 13 and 14 below.

We assume that $J = \{J(p,q); p \neq q \in S\}$ is a system of curves such that each $J(p,q)$ fulfills axiom A_1 of Definition 1.

Lemma 13 *With the above assumptions, we have*

$$\mathbb{R}^2 = \bigcup_{p \in S} \overline{VR(p,S)} \iff \forall p,q,r \in S: \quad D(p,q) \cap D(q,r) \subseteq D(p,r).$$

Proof. The direction from left to right has been shown in Lemma 5 without using axiom A_2. To prove the converse direction, let $z \in \mathbb{R}^2$. For an arbitrary $\epsilon > 0$ let $U := U_\epsilon(z)$ be an ϵ-neighborhood of z. As long as there exists a set $D = D(p,q)$ such that $U \cap D \neq \emptyset$ but $U \not\subset D$ we replace U by its open subset $U \cap D$. This process terminates after at most $\binom{n}{2}$ many steps. For the final set U, and for each pair $p \neq q$ of points from S we have

$$U \subset D(p,q) \;\text{ or }\; U \subset D(q,p). \tag{3}$$

The relation

$$p < q \;:\Leftrightarrow\; U \subset D(p,q)$$

is anti-symmetric, transitive by the right hand side of our lemma, and, by fact (3), either $p < q$ or $q < p$ must hold. Thus, $<$ defines a total order on S. Let p_ϵ denote the minimum element with respect to $<$ in S. Then, for each point $q \neq p_\epsilon$ in S we have $U \subset D(p_\epsilon, q)$ which implies $U \subset \mathrm{VR}(p_\epsilon, S)$. Thus, $U_\epsilon(z)$ contains points of $\mathrm{VR}(p_\epsilon, S)$. As ϵ tends to 0, the index p_ϵ may vary, but since S is finite there must be a subsequence of ϵ tending to 0 for which all p_ϵ are the same p. Consequently, $z \in \overline{\mathrm{VR}(p,S)}$. □

Lemma 14 *With the assumptions from above the following holds. If each set $VR(p,S')$, where $S' \subseteq S$ and $|S'| = 3$, is path-connected then each Voronoi region with respect to some $T \subseteq S$ is path-connected, too.*

Proof. If $T = \{p,q\}$ then $\mathrm{VR}(p,T) = D(p,q)$ is path-connected. For $|T| = 3$ the claim follows by assumption. Let $T \subseteq S$ be of size ≥ 4, and consider two points $x, y \in \mathrm{VR}(p,T)$. Let $t \neq t' \in T$ be different from p. By induction, there exist an arc π that connects x,y in $\mathrm{VR}(p, T \setminus \{t\})$, and a connecting arc π' in $\mathrm{VR}(p, T \setminus \{t'\})$. All points contained in both π and π' belong to $\mathrm{VR}(p,T)$. Suppose that some point $z \in \pi$ is not contained in $\mathrm{VR}(p,T)$, and let f, g be the first points of π' one meets when traversing π in both directions away from z. The two segments $\pi_{f,g}$ and $\pi'_{f,g}$ of π, π' between f and g form a domain, D; see Figure 9 (i).

We claim that there exists an arc $\alpha_{f,g} \subset D \cap \mathrm{VR}(p,T)$ from f to g. Then, by simultaneously replacing all (countably many) segments $\pi_{f,g}$ of π with $\alpha_{f,g}$, a path connecting x, y in $\mathrm{VR}(p,T)$ will result, thus proving our lemma.

This claim is shown as follows. Arcs $\pi_{f,g}$ and $\pi'_{f,g}$ together form a loop in $\mathrm{VR}(p, T \setminus \{t,t'\})$. This region is simply-connected by Lemma 12; one should observe that axioms A_2, A_3 have not been used in its proof. Hence, domain D too is subset of $\mathrm{VR}(p, T \setminus \{t,t'\})$. Now we distinguish two cases. First, we assume that the union

$$R := \mathrm{VR}(t,T) \;\cup\; \mathrm{VR}(t',T)$$

does not separate f from g in D. Then $D \setminus R$ belongs to $\mathrm{VR}(p,T)$ and contains the desired arc.

In the second case, R does separate f from g in D. We observe that $\mathrm{VR}(t',T)$ cannot intersect π, which is contained in $\mathrm{VR}(p, T \setminus \{t\})$, and $\mathrm{VR}(t,T)$ cannot intersect π'. Thus, both Voronoi regions have non-empty intersections with D. Now we consider the index set $S' := \{p,t,t'\}$. By assumption, there exists an arc ρ connecting f and g in $\mathrm{VR}(p,S')$. Arc ρ must avoid the union, R', of the Voronoi regions of t and t' with respect to S', which includes R. In the presence of $D \cap R'$, arc ρ can be homotopic with either $\pi_{f,g}$ or $\pi'_{f,g}$. In the first case, which is depicted in Figure 9 (ii), ρ and $\pi'_{f,g}$ together encircle points of $\mathrm{VR}(t,S') \subset D(t,p)$ although they are both contained in $D(p,t)$, which is simply-connected. In the second case, ρ and $\pi_{f,g}$ encircle points of $\mathrm{VR}(t',S') \subset D(t',p)$, but are themselves contained in $D(p,t')$—again a contradiction. $\square$

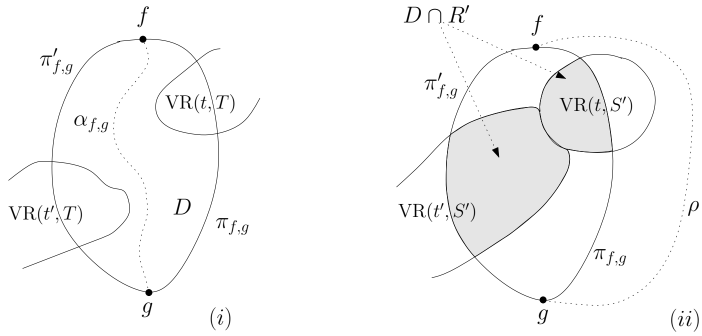

Figure 9: In (i), f and g can be connected by an arc in $\mathrm{VR}(p,T) \cap D$. In (ii), a path ρ connecting f and g in $\mathrm{VR}(p,\{p,t,t'\})$ must go around the Voronoi region of t or t'.

Now we can state the result of this subsection.

Theorem 15 *Let $J = \{J(p,q); p \neq q \in S\}$ be a system of curves each of which projects onto a closed Jordan curve through the north-pole of the sphere. Suppose that for each subset $S' \subseteq S$ of size 3 the Voronoi regions $VR(p,S')$ are path-connected, and that their closures cover the plane. Then J is admissible in the sense of Definition 1.*

5 Construction of AVDs

In this section we demonstrate that both algorithms developed for constructing abstract Voronoi diagrams, divide & conquer [21] and randomized incremental construction [24], also work for the AVD model presented in this paper. Since either algorithm was based on its own set of axioms, denoted by DC-AVD for divide & conquer and by RIC-AVD for randomized incremental construction, we also discuss how these sets differ from the new axioms introduced in Definition 1.

Since the bisecting curves our algorithms are dealing with can be rather complex objects, some care is required in establishing the cost of building an AVD. If all curves in J were algebraic of bounded degree, they could be described in constant space. Also, whatever elementary operations on bisecting curves are typically necessary, could be carried out in constant time. In general, even an elementary task, like testing two curves for intersection, may be undecidable. Therefore it seems reasonable to separate these issues from the task of constructing $V(S)$. To this end, we shall assume that a set of elementary operations on bisecting curves, that may depend on the algorithm considered, is capsuled in a basic module. Only this module can access the curves directly. Each call to the basic module performed by our algorithm is considered one step, just like a standard RAM operation. The algorithm is not charged for the basic module's real time and space consumption, which depends on the complexity of the bisecting curves $J(p,q)$.

5.1 Randomized incremental construction

A randomized incremental algorithm for abstract Voronoi diagrams was first presented by Mehlhorn et al. [27] and then generalized in [24]. The latter paper is based on the following assumptions.

As in Definition 1, a set S of n indices is given, and for each pair $p \neq q$ of indices in S a curve $J(p,q) = J(q,p)$ that splits the plane into two unbounded domains, $D(p,q)$ and $D(q,p)$. In addition, a total order $\prec$ on S is assumed to be part of the input. Now the *extended* Voronoi region of p with respect to S is defined by

$$\mathrm{EVR}(p,S) := \bigcap_{q \in S \setminus \{p\}} R(p,q),$$

where

$$R(p,q) := D(p,q)$$

if $q \prec p$, and

$$R(p,q) := D(p,q) \cup J(p,q),$$

if $p \prec q$. In addition, a *regular* Voronoi region was defined as

$$V(p,S) := \mathrm{EVR}^\circ(p,S),$$

where E° denotes the interior of E. Finally, the Voronoi diagram is defined as the union of the boundaries of all extended Voronoi regions $\mathrm{EVR}(p,S)$. Then the following properties are required.

Definition 16 *The pair $(J,\prec)$, where $J := \{J(p,q); p \neq q \in S\}$, is called* admissible for RIC-AVDs *if the following axioms are fulfilled.*

R_1) *Each curve $J(p,q)$, where $p,q \in S$, is mapped to a closed Jordan curve through the north pole by stereographic projection to the sphere.*

R_2) *For any p,q,r,s in S, the intersection $J(p,q) \cap J(r,s)$ consists of at most finitely many connected components.*

R_3) *For each subset $S' \subseteq S$ and for each $p' \in S'$, if $EVR(p',S')$ is non-empty then $V(p',S')$ is non-empty, too, and both sets are path-connected.*

R_4) *For each subset $S' \subseteq S$, we have*

$$\mathbb{R}^2 = \bigcup_{p' \in S'} EVR(p',S').$$

The following lemma shows a close connection between RIC-AVDs and our new AVD concept.

Lemma 17 *If, for some order $\prec$ on S, curve system J is admissible for RIC-AVDs then J is admissible for AVDs in the sense of Definition 1, and for each $S' \subseteq S$ and $p' \in S'$ we have $VR(p',S') = V(p',S')$.*

Proof. First, we prove the equality of Voronoi regions. Let $z \in V(p', S')$. By definition, z is an interior point of the extended region of p', so there is a neighborhood U of z contained in $\mathrm{EVR}(p', S')$. Clearly, U cannot intersect a bisecting curve $J(p', q')$; otherwise, U would contain points of $D(q', p')$, but such points do not belong to $\mathrm{EVR}(p', S')$. Therefore, $z \in U$ lies in all sets $D(p', q'), q' \in S'$, hence in $\mathrm{VR}(p', S')$. The converse inclusion is obvious.

Now assume that curve system J is admissible for RIC-AVDs. By the above, each Voronoi region $\mathrm{VR}(p', S') = V(p', S')$ is path-connected, thus proving axiom A_2. A point $z \in \mathbb{R}^2$ not contained in any region $\mathrm{VR}(s', S')$ belongs to some set $\mathrm{EVR}(p', S') \setminus V(p', S')$, by R_4. By the "piece of pie" Fact 1, p. 163 [24], z is contained in the closure of some region $V(q', S') = \mathrm{VR}(q', S')$, which proves A_3. Hence, J is admissible for AVDs. □

In Definition 2, p. 163 [24], Voronoi vertices and Voronoi edges were defined in the same way as for our Voronoi diagram, namely as (maximal) sets contained in the closures of two (resp.: of at least three) *regular* Voronoi regions.[12] By Lemma 17, regular Voronoi regions are the same as ours. Thus, the two definitions yield identical graphs as Voronoi diagrams.

The only difference lies in the fact that in RIC-AVDs, Voronoi edges are, as point sets, distributed among the Voronoi regions. For example, Figure 10 (i) shows a curve system admissible for each set of axioms considered in this paper. In (iii) the extended Voronoi region $\mathrm{EVR}(p, S)$ of p consists of the closure of its interior, $V(p, S)$, plus segment σ of the adjacent Voronoi edge. In fact, we have

$$\sigma \subseteq J(p, q) \cap J(p, r) \subset R(p, q) \cap R(p, r) = \mathrm{EVR}(p, S)$$

because of $p \prec q, r$. The rest, τ, of this edge belongs to the extended region of q since $q \prec r$. Point v is a Voronoi vertex of the RIC-AVD but the point w, where σ and τ meet, is not, because it lies in the closure of only two regular Voronoi regions.

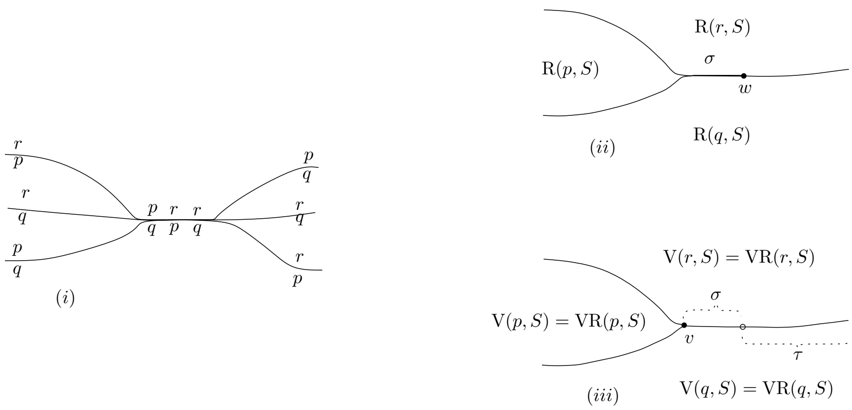

Figure 10: (i) A curve system for $S := \{p, q, r\}$ that is admissible under all definitions. (ii) Voronoi regions in the DC model, assuming $p \prec q \prec r$. Segment σ is contained in the region of p. (iii) The regular regions in the RIC model equal the Voronoi regions in our new model.

The RIC algorithm from [24] iteratively adds a random index s to the set R already considered, and updates both the Voronoi diagram $V(R)$ and a history graph, $H(R)$. This incremental step depends only on the structure of the intersection

$$\overline{V(s, R \cup \{s\})} \cap V(R);$$

[12] For the definition of Voronoi edges and vertices in our model, compare the end of the proof of our Theorem 10.

see the analysis starting on p. 165 [24]. This set, in turn, does not depend on the order $\prec$ on S. Therefore, we can run the incremental algorithm on a system of curves, admissible according to our Definition 1, without supplying an order $\prec$, and the Voronoi diagram $V(S)$ will be constructed. No harm can come from the fact that the finite intersection property R_2 is missing, for the following reasons. All geometric operations, in particular those on bisecting curves, are capsuled in a basic module, as was explained in the introduction to this section. With RIC-AVDs, this module takes as input a subset S' of five indices of S, and outputs a combinatorial description of $V(S')$; see p. 169 [24]. The incremental algorithm itself works in a purely combinatorial way on the outputs of this module. Its correctness does depend on the fact that $V(S)$ is a finite plane graph and on the piece of pie fact, but these are guaranteed by our Theorem 10 and by Lemma 11.

We have thus obtained the following counterpart of Theorem 2, p. 181 [24].

Theorem 18 *Let J be a curve system admissible in the sense of Definition 1 for index set S, where $|S| = n$. Then the abstract Voronoi diagram $V(S)$ can be constructed in an expected number of $O(n \log n)$ steps and in expected $O(n)$ space, by randomized incremental construction. A single step may involve a call to a basic module that returns a combinatorial description of a diagram of size five.*

Lemma 17 shows that curve systems admissible for RIC-AVDs form a subclass of curve systems admissible for new AVDs. This inclusion is strict, for two reasons. Obviously, bisecting curves are now allowed to intersect infinitely often. Also, abandoning the total order on S makes more curve systems admissible, as the following lemma shows.

Lemma 19 *There are admissible curve systems enjoying the finite intersection property, which are not admissible for RIC-AVDs under any order on index set S.*

Proof. Figure 11 shows a curve system admissible under Definition 1, with Voronoi edges drawn bold. For each permutation ijk of 123 there exists a point $z_{ijk} \in J(p_i, p_j) \cap J(p_j, p_k) \cap D(p_k, p_i)$. Under order $p_i \prec p_j \prec p_k$ we would have $z_{ijk} \in \mathrm{R}(p_i, p_j) \cap \mathrm{R}(p_j, p_k) \cap \mathrm{R}(p_k, p_i)$. This prevents z_{ijk} from being contained in *any* extended Voronoi region, thus violating condition R_4. $\square$

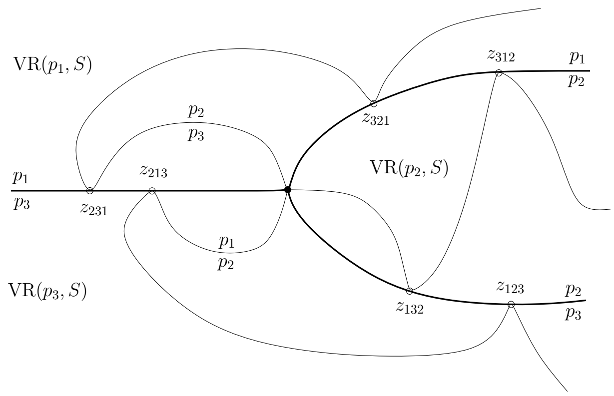

Figure 11: An admissible curve system that cannot be made admissible for RIC- or DC-AVDs by any order on S.

Summarizing, we have seen that under the new AVD axioms the incremental algorithm becomes both more natural and more powerful.

5.2 Divide & Conquer

The divide & conquer algorithm of [21] was designed for the original definition of abstract Voronoi diagrams that will now be reviewed. As in Subsection 5.1, the input curves $J(p,q)$ and the order $\prec$ on S were used to define Voronoi regions

$$R(p,S) = \mathrm{EVR}(p,S)$$

that are equal to the extended Voronoi regions for RIC-AVDs. Again, the Voronoi diagram $V(S)$ was defined as the union of all region boundaries. The important difference is that for DC-AVDs the interiors of the Voronoi regions were not required to be path-connected. More precisely, only the following properties were stated.

Definition 20 *The pair $(J, \prec)$, where $J := \{J(p,q); p \neq q \in S\}$, is called* admissible for DC-AVDs *if the following axioms are fulfilled.*

D_1) *Each curve $J(p,q)$, where $p,q \in S$, is mapped to a closed Jordan curve through the north pole by stereographic projection to the sphere.*

D_2) *For any p,q,r,s in S, the intersection $J(p,q) \cap J(r,s)$ consists of at most finitely many connected components.*

D_3) *For each subset $S' \subseteq S$ and for each $p' \in S'$, the set $R(p',S')$ is path-connected and has a non-empty interior.*

D_4) *For each subset $S' \subseteq S$, we have*

$$\mathbb{R}^2 = \bigcup_{p' \in S'} R(p',S').$$

Due to D_4, each point of the plane belonged to a Voronoi region. Since only the Voronoi regions, but not necessarily their interiors, need be path-connected, Voronoi regions could contain cut-points. A simple example is shown in Figure 12. Because of $p \prec q,r$, we have $J(p,q) \subset R(p,q)$ and $J(p,r) \subset R(p,r)$, so that segment $\sigma \subset J(p,q) \cap J(p,r)$ belongs to $R(p,q) \cap R(p,r) = R(p,S)$ where $S = \{p,q,r\}$. All points of segment σ are cut-points of the Voronoi region $R(p,S)$, whose removal disconnects this region.

In preparation for a precise definition of Voronoi edges, a segment σ of $J(p,q)$ was called a $\{p,q\}$-borderline if $p \prec q$ and $\sigma \subset R(p,S) \cap \overline{R(q,S)}$, or vice versa; compare Definition 2.3.4, p. 40 [21]. In Figure 12, σ is both, a $\{p,q\}$ and a $\{p,r\}$ borderline. In a way, the boundary of $R(p,S)$ has been squeezed together along σ. But the DC-AVD algorithm, that will be sketched below, operated as if the two borderlines forming σ were disjoint.

More complicated situations could arise at such points where more than two Voronoi regions met. Let us take a look at the pieces of pie around the point v depicted in Figure 13. Since Voronoi regions are connected, by D_3, and because their closures are simply-connected (see Lemma 12 and the discussion preceding it), only the Voronoi region to which point v belongs could contribute several pieces of pie to the neighborhood of v. These pieces were connected via v. In Figure 13, for example, point v belongs to $R(p,S)$ and connects the three pieces of this Voronoi region.

At point v, the Voronoi region of p was conceptually thickened, thus splitting v into as many "induced" points v_i as there are pieces of $R(p,S)$; compare Definition 2.5.1, p. 46 [21]. In Figure 13 this thickening results in three induced points, v_1, v_2, v_3. Of the borderlines adjacent to v, each v_i inherits those contained in the wedge between two consecutive p-pieces. Each

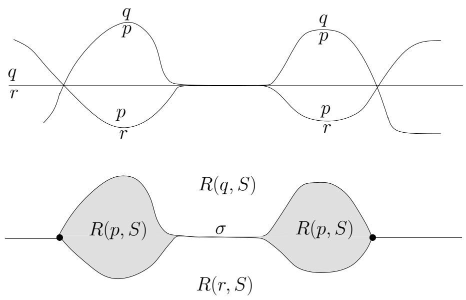

Figure 12: A curve system that satisfies the DC-AVD axioms, provided that order $\prec$ assigns both $J(p, q)$ and $J(p, r)$ to p, that is, $p \prec q, r$. The region $R(p, S)$ of p consists of the two shaded regions plus the segment σ connecting them.

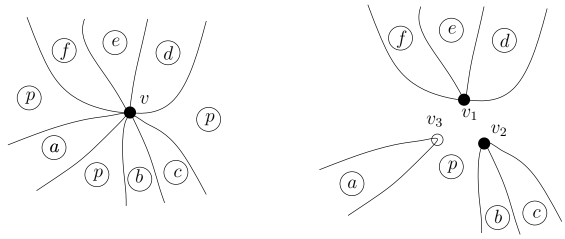

Figure 13: A pie neighborhood in a DC-AVD. Encircled indices denote Voronoi regions. Point v belongs to the region of $p \prec a, b, c, d, e, f$. It induces three points, v_1, v_2, v_3. Of those, v_1 and v_2 are Voronoi vertices of degree 4 resp. 3.

induced point adjacent to ≥ 3 borderlines was called a Voronoi vertex, and each borderline connecting two induced vertices was called a Voronoi edge. This definition determines the graph structure of DC-AVD. It differs from the way RIC-AVDs and new AVDs are defined as graphs. For example, in Figure 10 (ii) point w is adjacent to the borderlines of $\{p, q\}, \{p, r\}, \{q, r\}$ and, therefore, a Voronoi vertex. Point v is not a Voronoi vertex because it induces two points of degree only 2.

Now we briefly review the divide & conquer algorithm of [21] for DC-AVDs. The set S was divided into two subsets, L and R, of about the same size. Once their diagrams had been recursively computed they must be merged into $V(S)$. The merge step involved computing all $\{p, q\}$-borderlines of $V(S)$ where $p \in L$ and $q \in R$. These borderlines can form unbounded chains or cycles. Of each chain, a starting segment could be picked up at infinity by the technique of Chew and Drysdale [10]. For cycles, no such approach was known to work with AVDs. Therefore, it was required, in addition to properties D_1 to D_4, that L and R form an acyclic partition, that is, that for all subsets $L' \subseteq L, R' \subseteq R$, the borderlines separating L'-regions from R'-regions do not contain cycles.

Bisecting chains were traced through $V(L)$ and $V(R)$ simultaneously. Three phenomena needed special attention that cannot occur in the classical divide & conquer algorithm for Euclidean Voronoi diagrams [32].

First, the same point v of $V(S)$ could be incident to several L/R-separating borderlines. It was necessary to assign these borderlines to the vertices v_i induced by v in such a way that, in the graph $V(S)$, bisecting chains pass consistently through induced vertices. To this end, a function *site-of* was employed to determine if the Voronoi region containing v in $V(S)$ belongs to L or to

R. In the first case, any bisecting chain visiting v would be continued by the counterclockwise first L/R-borderline adjacent to v, while in the second case the clockwise first borderline around v was chosen; see pp. 93 and 112 in [21].

The *second* phenomenon concerned the tracing of a bisecting chain. Let $e \subset J(p,q)$ be an edge of chain K we are currently tracing through Voronoi region $R(p,L)$, where $p \in L$ is to the left of $J(p,q)$, as seen in tracing direction, and $q \in R$. Let v_L and v_R denote the points where the part of $J(p,q)$ extending e hits the boundaries of $R(p,L)$ and $R(q,R)$, respectively. If v_R lies before v_L on $J(p,q)$, edge e ends in v_R, and chain K is continued by an edge $e' \subset J(p,q')$, for some $q' \in R$.[13] Now we need to determine the point v'_L where $J(p,q')$ hits the boundary of $R(p,L)$. In the Euclidean case, v'_L can be efficiently located by scanning $\partial R(p,L)$ counterclockwise, starting from v_L. With AVDs this approach does not work, because $J(p,q')$ can cross $J(p,q)$ before reaching $\partial R(p,L)$; see Figure 3.19 on p. 84 [21].

This problem was solved in the following way. One maintains the last bisecting edge extension $T = J(p,q_s)$ whose endpoint on $\partial R(p,L)$ has actually been determined by counterclockwise scan. Now, in order to find the endpoint of $e' \subset J(p,q')$, one first tests if $J(p,q')$ crosses T. If so, edge e' ends there, and chain K is continued by a piece of T. Otherwise, we scan $\partial R(p,L)$ counterclockwise from v_L, as usual, find v'_L, and update T. This procedure was built into function *L-endpoint* on p. 85 of [21].

The *third* phenomenon was that an edge $e \subset J(p,q)$ does not necessarily end at the first point v of intersection of $J(p,q)$ with some other curve $J(p,p')$. In fact, if the two curves only touch at v then e may well continue beyond v. It does end when it comes to a *cross-point* where the status of $J(p,q)$ as a $\{p,q\}$-borderline ends; see p. 109 in [21].

As with RIC-AVDs, the time spent on constructing a DC-AVD was measured by the number of standard RAM operations plus the number of calls to some basic module, that contained the following elementary operations on bisecting curves.

- E_1) Given a curve $J(p,q)$ and a point v, determine if $v \in D(p,q)$ holds.
- E_2) Given a point $v \in J(p,q) \cap J(r,s) \cap J(t,u)$ and orientations of the curves, determine if $J(r,s)^+$ is prior to $J(t,u)^+$ in clockwise direction from $J(p,q)^+$, in a neifgborhood of v. Here, J^+ denotes the curve segment outgoing from v.
- E_3) Given points $v \in J(p,q)$, $w \in J(p,r)$ and orientations, determine the first cross-point on $J(p,r)$ after w caused by intersection with the part of $J(p,q)$ after v.
- E_4) Given a bisecting curve J with orientation, and points $v, w, x \in J$, determine if they appear in the order (v,w,x) on J.

Under these assumptions, the merge step could be completed in $O(n)$ many steps, where $n = |S|$; see Theorem 3.4.3.2 in [21].

Now we set out to prove the counterpart of this result for new AVDs.

Elementary operations E_1 and E_4 can remain unchanged. Operation E_2 can also be re-used if the predicate "prior to" is implemented correctly. It must include the situation where one curve touches the other in every neighborhood of v from the same side. In Figure 5, for example, F^- is prior to G^-, which is prior to F^+, which in turn is prior to G^+, in clockwise direction[14].

[13] For simplicity we assume that v_R is an interior point of the region $R(p,L)$, an interior edge point of $V(R)$, and not a multiple boundary point of $V(R)$.

[14] Observe that when we scan a vertex for an edge that continues an incoming bisecting chain, we need to evaluate the "prior to" predicate for only such pairs of bisecting curves that share one index. Lemma 6 rules out that one of them alternates infinitely often between the open halfplanes defined by the other.

Only operation E_3 must be modified. Instead of searching for a cross-point in the sense of DC-AVDs, we now determine the first Voronoi vertex. The modified operation is as follows.

E'_3) Given points $v \in J(p,q)$, $w \in J(p,r)$ and orientations, determine the first Voronoi vertex of $V(\{p,q,r\})$ that is situated on $J(p,q)$ after v, and on $J(p,r)$ after w, or return that no such vertex exists.

Now we show how to adapt the divide & conquer algorithm for DC-AVDs from [21] to the new AVD concept.

Theorem 21 *Let J be an admissible curve system for index set S of size n, and assume that $S = L \cup R$ is an acyclic partition. Then the Voronoi diagrams of $V(L)$ and $V(R)$ can be merged into $V(S)$ in $O(n)$ many steps and $O(n)$ space. A step may involve a standard RAM-Operation or one of the elementary bisector operations E_1, E_2, E'_3, E_4 stated above.*

Proof. It is crucial for the merge step that $V(L), V(R)$, and $V(S)$ are finite plane graphs. For DC-AVDs, this fact was a consequence of the finite intersection property D_2, and of a "piece of pie" lemma for arbitrary bisecting curves passing through a common point. For new AVDs, Theorem 10 ensures finiteness, and Lemma 11 establishes a piece of pie fact for the Voronoi edges emanating from a Voronoi vertex.

These facts allow us to run the DC-AVD merge algorithm, with the following two modifications. First, the algorithm can be greatly simplified. Multiple induced points, as depicted in Figure 13, can no longer occur, because each Voronoi region appears only once in any pie neighborhood of v, by Lemma 11. As a consequence, there is no need for different scan directions in computing bisecting chains. When edge e of chain K ends at vertex v, we simultaneaously scan the regions of $V(L)$ and $V(R)$ in clockwise direction around v for the first Voronoi edge e' that separates two regions from L and R, respectively. Then K is extended by e'.

The second modification is effected by replacing E_3 with E'_3. Its effect can be observed in Figure 10 (ii). Point w would be the cross-point where the borderline between the regions of p and q ends, but point v depicted in (iii) is the correct Voronoi vertex we need to compute in the new AVD model. This modification applies whenever two bisecting curves $J(p,q)$ and $J(p,r)$ are tested for "proper" intersection. □

The following consequence of Theorem 21 is immediate.

Theorem 22 *The abstract Voronoi diagram $V(S)$ can be constructed in $O(n \log n)$ many steps and in $O(n)$ space, if acyclic partitions can be found recursively in linear time. A single step may involve a call to a basic module that contains four elementary operations on bisecting curves.*

For convex distance functions, for example, standard split lines are guaranteed to give acyclic partitions. Further examples of acyclic partitions were provided in Chapter 4 [21].

Admissible curve systems under Definition 1 are not allowed to yield Voronoi regions whose interior is disconnected. Therefore, a curve system admissible for DC-AVDs need not be admissible for new AVDs. On the other hand, the curve system depicted in Figure 11 is admissible for new AVDs, but cannot be made admissible for DC-AVDs by any order; the same proof as for Lemma 19 applies. Thus, the divide & conquer algorithm for new AVDs becomes considerably simpler, while its scope has been shifted.

6 Conclusions

We have shown that the finite intersection property of bisecting curves is not necessary for defining and computing abstract Voronoi diagrams. With the simplified set of axioms suggested in this paper, the AVD concept has become more versatile as before. We expect that these axioms will also be helpful in further generalizing the AVD concept. Interesting open questions are how to deal with closed bisecting curves and with disconnected Voronoi regions. Also, further progress towards a general theory of 3-dimensional AVDs would be very valuable.

References

[1] M. Abellanas, F. Hurtado, and B. Palop. Transportation Networks and Voronoi Diagrams. Proceedings International Symposium on Voronoi Diagrams in Science and Engineering, 2004.

[2] A. Aggarwal, L. J. Guibas, J. Saxe, and P. W. Shor. A Linear Time Algorithm for Computing the Voronoi Diagram of a Convex Polygon. Discrete and Computational Geometry 4, pp. 591-604, 1989.

[3] H.-K. Ahn, O. Cheong, and R. v. Oostrum. Casting a Polyhedron with Directional Uncertainty. Computational Geometry: Theory and Applications 26(2), pp. 129–141, 2003.

[4] O. Aichholzer, F. Aurenhammer, and B. Palop. Quickest Paths, Straight Skeletons, and the City Voronoi Diagram. Discrete and Computational Geometry 31(7), pp. 17–35, 2004.

[5] M. McAllister, D. Kirkpatrick, and J. Snoeyink. A Compact Piecewise-Linear Voronoi Diagram for Convex Sites in the Plane. Discrete and Computational Geometry 15, pp. 73–105, 1996.

[6] F. Aurenhammer. Voronoi Diagrams: A Survey of a Fundamental Geometric Data Structure. ACM Computing Surveys 23(3), pp. 345–405, 1991.

[7] F. Aurenhammer and R. Klein. Voronoi Diagrams. In: J.R. Sack and G. Urrutia (Eds.), Handbook on Computational Geometry, Elsevier, pp. 201–290, 1999.

[8] S. W. Bae and K.-Y. Chwa. Voronoi Diagrams for a Transportation Network on the Euclidean Plane. International Journal on Computational Geometry and Applications 16, pp. 117-144, 2006.

[9] J. D. Boissonnat, C. Wormser, and M. Yvinec. Curved Voronoi Diagrams. In: J. D. Boissonnat and M. Teillaud (Eds.), Effective Computational Geometry for Curves and Surfaces. Springer, Mathematics and Visualization, 2006.

[10] P. Chew and S. Drysdale. Voronoi Diagrams Based on Convex Distance Functions. In: Proceedings 1st ACM Symposium on Computational Geometry, pp. 235–244, 1985.

[11] K. L. Clarkson and P. W. Shor. Applications of Random Sampling in Computational Geometry, II. Discrete and Computational Geometry 4, pp. 387–421, 1989.

[12] A. G. Corbalan, M. Mazon, and T. Recio. Geometry of Bisectors for Strictly Convex Distance Functions. International Journal of Computational Geometry and Applications 6(1), pp. 45-58, 1996.

[13] R. Descartes. Principia Philosophiae. Ludovicus Elzevirius, Amsterdam, 1644.

[14] H. Edelsbrunner and R. Seidel. Voronoi Diagrams and Arrangements. Discrete and Computational Geometry 1, pp. 387–421, 1986.

[15] S. Fortune. Voronoi Diagrams and Delaunay Triangulations. In: J.E. Goodman and J. O'Rourke (Eds.), Handbook of Discrete and Computational Geometry, Chapter 20, CRC Press LLC, pp. 377–388, 1997.

[16] P. Gruber. Convex and Discrete Geometry. Grundlehren der mathemetischen Wissenschaften, Vol. 336, Springer-Verlag, 2007.

[17] M. Held. On the Computational Geometry of Pocket Machining. Lecture Notes in Computer Science 500, Springer-Verlag, 1991.

[18] Ch. Icking, R. Klein, N.-M. Le, and L. Ma. Convex Distance Functions in 3-Space are Different. Fundamenta Informaticae 22(4), pp. 331–352, 1995.

[19] M. I. Karavelas and M. Yvinec. The Voronoi Diagram of Planar Convex Objects. 11th European Symposium on Algorithms (ESA'03). LNCS 2832, pp. 337–348, 2003.

[20] R. Klein. Abstract Voronoi Diagrams and Their Applications. In: H. Noltemeier (Ed.), Proceedings Workshop on Computational Geometry and its Applications, LNCS 333, pp. 148–157, 1988.

[21] R. Klein. Concrete and Abstract Voronoi Diagrams. Lecture Notes in Computer Science 400, Springer-Verlag, 1987.

[22] R. Klein. Combinatorial Properties of Abstract Voronoi Diagrams. In: M. Nagl (Ed.), Graph-Theoretic Concepts in Computer Science (WG '89), Rolduc, Netherlands. Lecture Notes in Computer Science 411, pp. 356–369, Springer-Verlag, 1989.

[23] R. Klein and A. Lingas. Hamiltonian Abstract Voronoi Diagrams in Linear Time. In: D.-Z. Du and X.-S. Zhang (Eds.), Algorithms and Computations, 5th International Symposium (ISAAC '94), LNCS 834, pp. 11–19, 1994.

[24] R. Klein, K. Mehlhorn, and St. Meiser. Randomized Incremental Construction of Abstract Voronoi Diagrams. Computational Geometry: Theory and Applications 3, pp. 157–184, 1993.

[25] N.-M. Lê. Abstract Voronoi Diagram in 3-Space. Journal of Computer and System Sciences 68(1), pp. 41–79, 2004.

[26] K. K. Malinauskas. Dynamic Construction of Abstract Voronoi Diagrams. Fundamentalnaya i Prikladnaya Matematika 13(2), pp. 133-146, 2007.

[27] K. Mehlhorn, St. Meiser, and C. Ó'Dúnlaing. On the Construction of Abstract Voronoi Diagrams. Discrete and Computational Geometry 6, pp. 211–224, 1991.

[28] K. Mehlhorn, St. Meiser, and R. Rasch. Furthest Site Abstract Voronoi Diagrams. International Journal of Computational Geometry and Applications 11(6), pp. 583–616, 2001.

[29] K. Menger. Untersuchungen über allgemeine Metrik I, II, III. Mathematische Annalen 100, pp. 75–163, 1928.

[30] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. Spatial Tessellations: Concepts and Applications of Voronoi Diagrams. Wiley Series in Probability and Statistics, 2000.

[31] W. Rinow. Topologie. VEB Deutscher Verlag der Wissenschaften, Leipzig 1975.

[32] M. I. Shamos and D. Hoey. Closest-Point Problems. In Proceedings 16th IEEE Symposium on Foundations of Computer Science, pp. 151–162, 1975.

[33] M. Seel. LEDA Extension Package *Abstract Voronoi Diagrams*. www.algorithmic-solutions.com/leda/leps/index.htm.

[34] C. Thomassen. The Converse of the Jordan Curve Theorem and a Characterization of Planar Maps. Geometriae Dedicata 32, pp. 53–57, 1989.