Deletion in Abstract Voronoi Diagrams in Expected Linear Time

Kolja Junginger
Faculty of Informatics, USI Università della Svizzera italiana,
Lugano, Switzerland
kolja.junginger@usi.ch

Evanthia Papadopoulou
Faculty of Informatics, USI Università della Svizzera italiana,
Lugano, Switzerland
evanthia.papadopoulou@usi.ch

Abstract

Updating an abstract Voronoi diagram in linear time, after deletion of one site, has been an open problem for a long time. Similarly for various concrete Voronoi diagrams of generalized sites, other than points. In this paper we present a simple, expected linear-time algorithm to update an abstract Voronoi diagram after deletion. We introduce the concept of a *Voronoi-like diagram*, a relaxed version of a Voronoi construct that has a structure similar to an abstract Voronoi diagram, without however being one. Voronoi-like diagrams serve as intermediate structures, which are considerably simpler to compute, thus, making an expected linear-time construction possible. We formalize the concept and prove that it is robust under an insertion operation, thus, enabling its use in incremental constructions.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Abstract Voronoi diagram, linear-time algorithm, update after deletion, randomized incremental algorithm

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.50

Related Version A full version of this paper is available at `http://arxiv.org/abs/1803.05372`.

Funding Supported in part by the Swiss National Science Foundation, DACH project SNF-200021E-154387

1 Introduction

The Voronoi diagram of a set S of n simple geometric objects, called sites, is a well-known geometric partitioning structure that reveals proximity information for the input sites. Classic variants include the *nearest-neighbor*, the *farthest-site*, and the *order-k* Voronoi diagram of S ($1 \leq k < n$). *Abstract Voronoi diagrams* [11] offer a unifying framework for various concrete and well-known instances. Some classic Voronoi diagrams have been well investigated, with optimal construction algorithms available in many cases, see e.g., [2] for references and more information or [16] for numerous applications.

For certain *tree-like* Voronoi diagrams in the plane, linear-time construction algorithms have been well-known to exist, see e.g., [1, 7, 13, 8]. The first technique was introduced by Aggarwal et al. [1] for the Voronoi diagram of points in convex position, given the order of points along their convex hull. It can be used to derive linear-time algorithms for other fundamental problems: (1) updating a Voronoi diagram of points after deletion of one site in

34th International Symposium on Computational Geometry (SoCG 2018).
Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 50; pp. 50:1–50:14
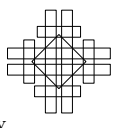

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

time linear to the number of the Voronoi neighbors of the deleted site; (2) computing the farthest Voronoi diagram of point-sites in linear time, after computing their convex hull; (3) computing the order-$(k+1)$ subdivision within an order-k Voronoi region. There is also a much simpler randomized approach for the same problems introduced by Chew [7]. Klein and Lingas [13] adapted the linear-time framework [1] to abstract Voronoi diagrams, under restrictions, showing that a *Hamiltonian abstract Voronoi diagram* can be computed in linear time, given the order of Voronoi regions along an unbounded simple curve, which visits each region *exactly once* and can intersect each bisector only *once*. This construction has been extended recently to include forest structures [4] under similar conditions, where no region can have multiple faces within the domain enclosed by a curve. The medial axis of a simple polygon is another well-known problem to admit a linear-time construction, shown by Chin et al. [8].

In this paper we consider the fundamental problem of updating a two-dimensional Voronoi diagram, after deletion of one site, and provide an expected linear-time algorithm to achieve this task. We consider the framework of abstract Voronoi diagrams to simultaneously address the various concrete instances under their umbrella. To the best of our knowledge, no linear-time construction algorithms are known for concrete diagrams of non-point sites, nor for abstract Voronoi diagrams. Related is our expected linear-time algorithm for the concrete farthest-segment Voronoi diagram [10][1], however, definitions are geometric, relying on star-shapeness and visibility properties of segment Voronoi regions, which do not extend to the abstract model. In this paper we consider a new formulation.

Abstract Voronoi diagrams. Abstract Voronoi diagrams (AVDs) were introduced by Klein [11]. Instead of sites and distance measures, they are defined in terms of bisecting curves that satisfy some simple combinatorial properties. Given a set S of n abstract sites, the bisector $J(p,q)$ of two sites $p,q \in S$ is an unbounded Jordan curve, homeomorphic to a line, that divides the plane into two open domains: the *dominance region of* p, $D(p,q)$ (having label p), and the *dominance region of* q, $D(q,p)$ (having label q), see Figure 1. The *Voronoi region* of p is

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

The *(nearest-neighbor) abstract Voronoi diagram* of S is $\mathcal{V}(S) = \mathbb{R}^2 \setminus \bigcup_{p \in S} \mathrm{VR}(p,S)$.

Following the traditional model of abstract Voronoi diagrams (see e.g. [11, 3, 6, 5]) the system of bisectors is assumed to satisfy the following axioms, for every subset $S' \subseteq S$:

- **(A1)** Each nearest Voronoi region $\mathrm{VR}(p,S')$ is non-empty and pathwise connected.
- **(A2)** Each point in the plane belongs to the closure of a nearest Voronoi region $\mathrm{VR}(p,S')$.
- **(A3)** After stereographic projection to the sphere, each bisector can be completed to a Jordan curve through the north pole.
- **(A4)** Any two bisectors $J(p,q)$ and $J(r,t)$ intersect transversally and in a finite number of points. (It is possible to relax this axiom, see [12]).

$\mathcal{V}(S)$ is a plane graph of structural complexity $O(n)$ and its regions are simply-connected. It can be computed in time $O(n \log n)$, randomized [14] or deterministic [11]. To update $\mathcal{V}(S)$, after deleting one site $s \in S$, we compute $\mathcal{V}(S \setminus \{s\})$ within $\mathrm{VR}(s,S)$. The sequence of

[1] A preliminary version contains a gap when considering the linear-time framework of [1], thus, a linear-time construction for the farthest segment Voronoi diagram remains an open problem.

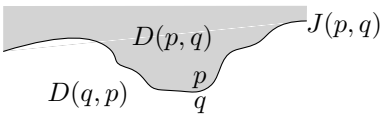

Figure 1 A bisector $J(p,q)$ and its dominance regions; $D(p,q)$ is shown shaded.

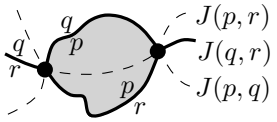

Figure 2 The Voronoi diagram $\mathcal{V}(\{p,q,r\})$ in solid lines. The shaded region is $\mathrm{VR}(p,\{p,q,r\})$.

site-occurrences along $\partial\mathrm{VR}(s,S)$ forms a Davenport-Schinzel sequence of order 2 and this constitutes a major difference from the respective problem for points, where no repetition can occur. $\mathcal{V}(S\setminus\{s\})\cap\mathrm{VR}(s,S)$ contains disconnected Voronoi regions, which introduce several complications. For example, $\mathcal{V}(S')\cap\mathrm{VR}(s,S'\cup\{s\})$ for $S'\subset S\setminus\{s\}$ may contain various faces that are not related to $\mathcal{V}(S\setminus\{s\})\cap\mathrm{VR}(s,S)$, and conversely, an arbitrary sub-sequence of $\partial\mathrm{VR}(s,S)$ need not correspond to any Voronoi diagram. At first sight, a linear-time algorithm may seem infeasible.

Our results. In this paper we give a simple randomized algorithm to compute $\mathcal{V}(S\setminus\{s\})$ within $\mathrm{VR}(s,S)$ in expected time linear on the complexity of $\partial\mathrm{VR}(s,S)$. The algorithm is simple, not more complicated than its counterpart for points [7], and this is achieved by computing simplified intermediate structures that are interesting in their own right. These are *Voronoi-like* diagrams, having a structure similar to an abstract Voronoi diagram, however, they are not Voronoi structures. *Voronoi-like regions* are supersets of real Voronoi regions, and their boundaries correspond to *monotone paths* in the relevant system of bisectors, rather than to an *envelope* in the same system as in a real Voronoi diagram (see Definition 5). We prove that Voronoi-like diagrams are well-defined, and also they are robust under an *insertion operation*, thus, making possible a randomized incremental construction for $\mathcal{V}(S\setminus\{s\})\cap\mathrm{VR}(s,S)$ in linear time. We expect the concept to find uses in other Voronoi computations, where computing intermediate relaxed structures may simplify the entire computation. A first candidate in this direction is the linear-time framework of Aggarwal et al. [1] that we plan to investigate next.

Our approach can be adapted (in fact, simplified) to compute in expected linear time the farthest abstract Voronoi diagram, after the sequence of its faces at infinity is known. The latter sequence can be computed in time $O(n\log n)$. We also expect that our algorithm can be adapted to compute the order-$(k{+}1)$ subdivision within an order-k abstract Voronoi region in expected time linear on the complexity of the region boundary.[2] Our technique can be applied to concrete diagrams that may not strictly fall under the AVD model such as Voronoi diagrams of line segments that may intersect and of planar straight-line graphs (including simple and non-simple polygons).

2 Preliminaries

Let S be a set of n abstract *sites* (a set of indices) that define an *admissible* system of bisectors in the plane $\mathcal{J}=\{J(p,q):p\neq q\in S\}$, which fulfills axioms (A1)–(A4) for every $S'\subseteq S$. The (nearest) Voronoi region of p is $\mathrm{VR}(p,S)=\bigcap_{q\in S\setminus\{p\}}D(p,q)$ and the Voronoi diagram of S is $\mathcal{V}(S)=\mathbb{R}^2\setminus\bigcup_{p\in S}\mathrm{VR}(p,S)$, see, e.g., Figure 2.

[2] The adaptation is non-trivial, thus, we only make a conjecture here and plan to consider details in subsequent work.

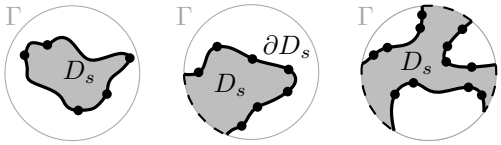

Figure 3 The domain $D_s = \mathrm{VR}(s, S) \cap D_\Gamma$.

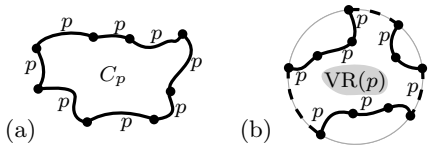

Figure 4 (a) A p-inverse cycle. (b) A p-cycle.

Bisectors that have a site p in common are called *p-related* or simply *related*; related bisectors can intersect at most twice [11, Lemma 3.5.2.5]. When two related bisectors $J(p, q)$ and $J(p, r)$ intersect, bisector $J(q, r)$ also intersects with them at the same point(s) [11], and these points are the Voronoi vertices of $\mathcal{V}(\{p, q, r\})$, see Figure 2. Since any two related bisectors in $\mathcal{J}$ intersect at most twice, the sequence of site occurrences along $\partial\mathrm{VR}(p, S)$, $p \in S$, forms a Davenport-Schinzel sequence of order 2 (by [19, Theorem 5.7]).

To update $\mathcal{V}(S)$ after deleting one site $s \in S$, we compute $\mathcal{V}(S \setminus \{s\})$ within $\mathrm{VR}(s, S)$, i.e., compute $\mathcal{V}(S \setminus \{s\}) \cap \mathrm{VR}(s, S)$. Its structure is given in the following lemma. Figure 7(a) illustrates $\mathcal{V}(S \setminus \{s\}) \cap \mathrm{VR}(s, S)$ (in red) for a bounded region $\mathrm{VR}(s, S)$, where the region's boundary is shown in bold.

▶ **Lemma 1.** *$\mathcal{V}(S \setminus \{s\}) \cap VR(s, S)$ is a forest having exactly one face for each Voronoi edge of $\partial VR(s, S)$. Its leaves are the Voronoi vertices of $\partial VR(s, S)$, and points at infinity if $VR(s, S)$ is unbounded. If $VR(s, S)$ is bounded then $\mathcal{V}(S \setminus \{s\}) \cap VR(s, S)$ is a tree.*

Let Γ be a closed Jordan curve in the plane large enough to enclose all the intersections of bisectors in $\mathcal{J}$, and such that each bisector crosses Γ exactly twice and transversally. Without loss of generality, we restrict all computations within Γ.[3] The curve Γ can be interpreted as $J(p, s_\infty)$, for all $p \in S$, where s_∞ is an additional site at infinity. Let the interior of Γ be denoted as D_Γ. Our *domain of computation* is $D_s = \mathrm{VR}(s, S) \cap D_\Gamma$, see Figure 3; we compute $\mathcal{V}(S \setminus \{s\}) \cap D_s$.

The following lemmas are used as tools in our proofs. Let C_p be a cycle of p-related bisectors in the arrangement of bisectors $\mathcal{J} \cup \Gamma$. If for every edge in C_p the label p appears on the outside of the cycle then C_p is called *p-inverse*, see Figure 4(a). If the label p appears only inside C_p then C_p is called a *p-cycle*, see Figure 4(b). By definition, $\mathrm{VR}(p, S) \subseteq C_p$ for any p-cycle C_p. A p-inverse cycle cannot contain pieces of Γ.

▶ **Lemma 2.** *In an admissible bisector system there is no p-inverse cycle.*

Proof. The farthest Voronoi region of p is $\mathrm{FVR}(p, S) = \bigcap_{q \in S \setminus \{p\}} D(q, p)$. By its definition, $\mathrm{FVR}(p, S)$ must be enclosed in any p-inverse cycle C_p. But farthest Voronoi regions must be unbounded [15, 3] deriving a contradiction. ◀

The following *transitivity lemma* is a consequence of transitivity of dominance regions [3, Lemma 2] and the fact that bisectors $J(p, q), J(q, r), J(p, r)$ intersect at the same point(s).

▶ **Lemma 3.** *Let $z \in \mathbb{R}^2$ and $p, q, r \in S$. If $z \in D(p, q)$ and $z \in \overline{D(q, r)}$, then $z \in D(p, r)$.*

We make a general position assumption that no three p-related bisectors intersect at the same point. This implies that Voronoi vertices have degree 3.

[3] The presence of Γ is conceptual and its exact position unknown; we never compute coordinates on Γ.

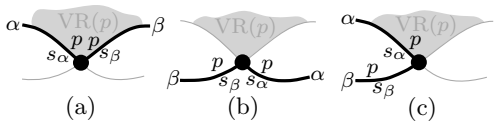

Figure 5 (a) Arcs α, β fulfill the p-monotone path condition; they do not fulfill it (b) and (c).

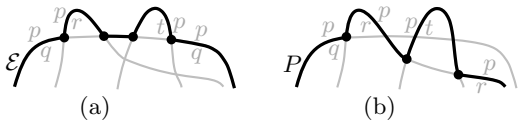

Figure 6 (a) The envelope $\mathcal{E} = \text{env}(\mathcal{J}_{p,\{q,r,t\}})$. (b) A p-monotone path P in $\mathcal{J}_{p,\{q,r,t\}}$.

3 Problem formulation and definitions

Let $\mathcal{S}$ denote the sequence of Voronoi edges along $\partial\text{VR}(s, S)$, i.e., $\mathcal{S} = \partial\text{VR}(s, S) \cap D_\Gamma$. We consider $\mathcal{S}$ as a cyclically ordered set of *arcs*, where each arc is a Voronoi edge of $\partial\text{VR}(s, S)$. Each arc $\alpha \in \mathcal{S}$ is induced by a site $s_\alpha \in S \setminus \{s\}$ such that $\alpha \subseteq J(s, s_\alpha)$. A site p may induce several arcs on $\mathcal{S}$; recall, that the sequence of site occurrences along $\partial\text{VR}(s, S)$ is a Davenport-Schinzel sequence of order 2.

We can interpret the arcs in $\mathcal{S}$ as sites that induce a Voronoi diagram $\mathcal{V}(\mathcal{S})$, where $\mathcal{V}(\mathcal{S}) = \mathcal{V}(S \setminus \{s\}) \cap D_s$ and $D_s = \text{VR}(s, S) \cap D_\Gamma$. Figure 7(a) illustrates $\mathcal{S}$ and $\mathcal{V}(\mathcal{S})$ in black (bold) and red, respectively. By Lemma 1, each face of $\mathcal{V}(S \setminus \{s\}) \cap D_s$ is incident to exactly one arc in $\mathcal{S}$. In this respect, each arc α in $\mathcal{S}$ has a Voronoi region, $\text{VR}(\alpha, \mathcal{S})$, which is the face of $\mathcal{V}(S \setminus \{s\}) \cap D_s$ incident to α.

For a site $p \in S$ and $S' \subseteq S$, let $\mathcal{J}_{p,S'} = \{J(p, q) \,|\, q \in S', q \neq p\}$ denote the set of all p-related bisectors involving sites in S'. The arrangement of a bisector set J is denoted by $\mathcal{A}(J)$. $\mathcal{A}(\mathcal{J}_{p,S'})$ may consist of more than one connected components.

▶ **Definition 4.** A *path* P in $\mathcal{J}_{p,S'}$ is a connected sequence of alternating edges and vertices of the arrangement $\mathcal{A}(\mathcal{J}_{p,S'})$. An *arc* α of P is a maximally connected set of consecutive edges and vertices of the arrangement along P, which belong to the same bisector. The common endpoint of two consecutive arcs of P is a *vertex of* P. An arc of P is also called an *edge*.

Two consecutive arcs in a path P are pieces of different bisectors. We use the notation $\alpha \in P$ for referring to an arc α of P. For $\alpha \in P$, let $s_\alpha \in S$ denote the site in S that *induces* α, where $\alpha \subseteq J(p, s_\alpha)$.

▶ **Definition 5.** A path P in $\mathcal{J}_{p,S'}$ is called p-*monotone* if any two consecutive arcs $\alpha, \beta \in P$, where $\alpha \subseteq J(p, s_\alpha)$ and $\beta \subseteq J(p, s_\beta)$, induce the Voronoi edges of $\partial\text{VR}(p, \{p, s_\alpha, s_\beta\})$, which are incident to the common endpoint of α, β (see Figure 5).

▶ **Definition 6.** The *envelope of* $\mathcal{J}_{p,S'}$, with respect to site p, is $\text{env}(\mathcal{J}_{p,S'}) = \partial\text{VR}(p, S' \cup \{p\})$, called a p-envelope (see Figure 6(a)).

Figure 6 illustrates two p-monotone paths, where the path in Figure 6(a) is a p-envelope. Notice, $\mathcal{S}$ is the envelope of the s-related bisectors in $\mathcal{J}$, $\mathcal{S} = \text{env}(\mathcal{J}_{s,S\setminus\{s\}}) \cap D_\Gamma$. A p-monotone path that is not a p-envelope can be a Davenport-Schinzel sequence of order > 2, with respect to site occurrences in $S \setminus \{s\}$.

The system of bisectors $\mathcal{J}_{p,S'}$ may consist of several connected components. For convenience, in order to unify the various connected components of $\mathcal{A}(\mathcal{J}_{p,S'})$ and to consider its p-monotone paths as single curves, we include the curve Γ in the corresponding system of bisectors. Then, $\text{env}(\mathcal{J}_{p,S'} \cup \Gamma)$ is a closed p-monotone path, whose connected components in $\mathcal{J}_{p,S'}$ are interleaved with arcs of Γ.

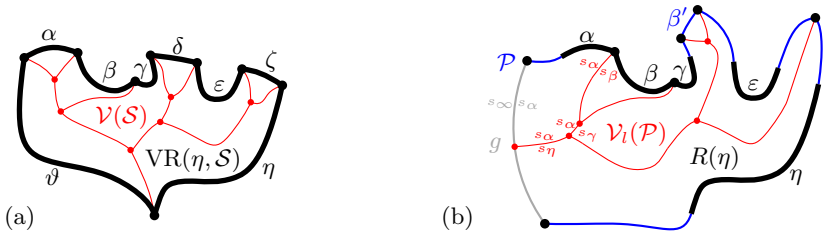

Figure 7 (a) illustrates $\mathcal{S}$ in black (bold) and $\mathcal{V}(\mathcal{S})$ in red, $\mathcal{S} = (\alpha, \beta, \gamma, \delta, \varepsilon, \zeta, \eta, \vartheta)$. (b) illustrates $\mathcal{V}_l(\mathcal{P})$ for a boundary curve $\mathcal{P} = (\alpha, \beta, \gamma, \beta', \varepsilon, \eta, g)$ for $\mathcal{S}'$, where $\mathcal{S}' = (\alpha, \beta, \gamma, \varepsilon, \eta)$ is shown in bold. The arcs of $\mathcal{P}$ are original except the auxiliary arc β' and the Γ-arc g.

▶ **Definition 7.** Consider $\mathcal{S}' \subseteq \mathcal{S}$ and let $S' = \{s_\alpha \in S \mid \alpha \in \mathcal{S}'\} \subseteq S \backslash \{s\}$ be its corresponding set of sites. A closed s-monotone path in $\mathcal{J}_{s,S'} \cup \Gamma$ that contains all arcs in $\mathcal{S}'$ is called a *boundary curve for* $\mathcal{S}'$. The part of the plane enclosed in a boundary curve $\mathcal{P}$ is called the *domain* of $\mathcal{P}$, and it is denoted by $D_\mathcal{P}$. Given $\mathcal{P}$, we also use notation $S_\mathcal{P}$ to denote S'.

A set of arcs $\mathcal{S}' \subset \mathcal{S}$ can admit several different boundary curves. One such boundary curve is its envelope $\mathcal{E} = \mathrm{env}(\mathcal{J}_{s,S'} \cup \Gamma)$. Figure 7(b) illustrates a boundary curve for $\mathcal{S}' \subseteq \mathcal{S}$, where $\mathcal{S}$ is the set of arcs in Figure 7(a).

A boundary curve $\mathcal{P}$ in $\mathcal{J}_{s,S'} \cup \Gamma$ consists of pieces of bisectors in $\mathcal{J}_{s,S'}$, called *boundary arcs*, and pieces of Γ, called Γ-*arcs*. Γ-arcs correspond to openings of the domain $D_\mathcal{P}$ to infinity. Among the boundary arcs, those that contain an arc of $\mathcal{S}'$ are called *original* and others are called *auxiliary arcs*. Original boundary arcs are expanded versions of the arcs in $\mathcal{S}'$. To distinguish between them, we call the elements of $\mathcal{S}$ *core arcs* and use an * in their notation. In Figure 7 the core arcs are illustrated in bold.

For a set of arcs $\mathcal{S}' \subseteq \mathcal{S}$, we define the Voronoi diagram of $\mathcal{S}' \subseteq \mathcal{S}$ as $\mathcal{V}(\mathcal{S}') = \mathcal{V}(S') \cap D_\mathcal{E}$, where $\mathcal{E}$ is the s-envelope $\mathrm{env}(\mathcal{J}_{s,S'} \cup \Gamma)$. $\mathcal{V}(\mathcal{S}')$ can be regarded as the Voronoi diagram of the envelope $\mathcal{E}$, thus, it can also be denoted as $\mathcal{V}(\mathcal{E})$. The face of $\mathcal{V}(\mathcal{S}')$ incident to an arc $\alpha \in \mathcal{E}$ is called the Voronoi region of α and is denoted by $\mathrm{VR}(\alpha, \mathcal{S}')$. We would like to extend the definition of $\mathcal{V}(\mathcal{S}')$ to any boundary curve stemming out of $\mathcal{S}'$. To this goal we define a *Voronoi-like diagram* for any boundary curve $\mathcal{P}$ of $\mathcal{S}'$. Notice, $D_s \subseteq D_\mathcal{E} \subseteq D_\mathcal{P}$.

▶ **Definition 8.** Given a boundary curve $\mathcal{P}$ in $\mathcal{J}_{s,S'} \cup \Gamma$, a *Voronoi-like diagram* of $\mathcal{P}$ is a plane graph on $\mathcal{J}(S') = \{J(p,q) \in \mathcal{J} \mid p, q \in S'\}$ inducing a subdivision on the domain $D_\mathcal{P}$ as follows (see Figure 7(b)):

1. There is exactly one face $R(\alpha)$ for each boundary arc α of $\mathcal{P}$, and $\partial R(\alpha)$ consists of the arc α plus an s_α-monotone path in $\mathcal{J}_{s_\alpha,S'} \cup \Gamma$.
2. $\bigcup_{\alpha \in \mathcal{P} \backslash \Gamma} \overline{R(\alpha)} = \overline{D_\mathcal{P}}$.

The Voronoi-like diagram of $\mathcal{P}$ is $\mathcal{V}_l(\mathcal{P}) = D_\mathcal{P} \setminus \bigcup_{\alpha \in \mathcal{P}} R(\alpha)$.

Voronoi-like regions in $\mathcal{V}_l(\mathcal{P})$ are related to real Voronoi regions in $\mathcal{V}(\mathcal{S}')$ as supersets, as shown in the following lemma. In Figure 7(b) the Voronoi-like region $R(\eta)$ is a superset of its corresponding Voronoi region $\mathrm{VR}(\eta, \mathcal{S})$ in (a); similarly for e.g., $R(\alpha)$. Note that not every boundary curve of $\mathcal{S}' \subset \mathcal{S}$ needs to admit a Voronoi-like diagram.

▶ **Lemma 9.** *Let α be a boundary arc in a boundary curve $\mathcal{P}$ of $\mathcal{S}'$ such that a portion $\tilde{\alpha} \subseteq \alpha$ appears on the s-envelope $\mathcal{E}$ of $\mathcal{S}'$, $\mathcal{E} = env(\mathcal{J}_{s,S'} \cup \Gamma)$. Given $\mathcal{V}_l(\mathcal{P})$, $R(\alpha) \supseteq VR(\tilde{\alpha}, \mathcal{S}')$. If α is original, then $R(\alpha) \supseteq VR(\tilde{\alpha}, \mathcal{S}') \supseteq VR(\alpha^*, \mathcal{S})$.*

Proof. By the definition of a Voronoi region, no piece of a bisector $J(s_\alpha, \cdot)$ can appear in the interior of $\mathrm{VR}(\tilde{\alpha}, \mathcal{S}')$, where $\tilde{\alpha} \in \mathcal{E}$ (recall that $\mathcal{V}(\mathcal{S}') = \mathcal{V}(\mathcal{E})$). Since in addition $\alpha \supseteq \tilde{\alpha}$, the

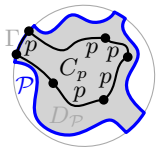

Figure 8 A p-cycle (possibly with Γ-arcs) within $\overline{D_{\mathcal{P}}}$ does not exist.

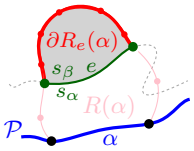

Figure 9 The shaded region lies in $D(s_\beta, s_\alpha)$.

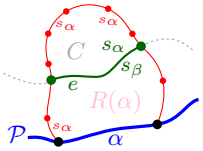

Figure 10 Edge e has a contradictory edge labeling.

claim follows. For an original arc α, since $S' \subseteq S$, by the monotonicity property of Voronoi regions, we also have $\mathrm{VR}(\tilde{\alpha}, \mathcal{S}') \supseteq \mathrm{VR}(\alpha^*, \mathcal{S})$. ◀

As a corollary to Lemma 9, the adjacencies of the real Voronoi diagram $\mathcal{V}(\mathcal{S}')$ are preserved in $\mathcal{V}_l(\mathcal{P})$, for all arcs that are common to the envelope $\mathcal{E}$ and the boundary curve $\mathcal{P}$. In addition, $\mathcal{V}_l(\mathcal{E})$ coincides with the real Voronoi diagram $\mathcal{V}(\mathcal{S}')$.

▶ **Corollary 10.** *$\mathcal{V}_l(\mathcal{E}) = \mathcal{V}(\mathcal{S}')$. This also implies $\mathcal{V}_l(\mathcal{S}) = \mathcal{V}(\mathcal{S})$.*

The following Lemma 12 gives a basic property of Voronoi-like regions that is essential for subsequent proofs. To establish it we first need the following observation.

▶ **Lemma 11.** *$\overline{D_{\mathcal{P}}}$ cannot contain a p-cycle of $\mathcal{J}(S_{\mathcal{P}}) \cup \Gamma$, for any $p \in S_{\mathcal{P}}$.*

Proof. Let $p \in S_{\mathcal{P}}$ define an original arc along $\mathcal{P}$. This arc is bounding $\mathrm{VR}(p, S_{\mathcal{P}} \cup \{s\})$, thus, it must have a portion within $\mathrm{VR}(p, S_{\mathcal{P}})$. Hence, $\mathrm{VR}(p, S_{\mathcal{P}})$ has a non-empty intersection with $\mathbb{R}^2 \setminus \overline{D_{\mathcal{P}}}$. But $\mathrm{VR}(p, S_{\mathcal{P}})$ must be enclosed within any p-cycle of $\mathcal{J}(S_{\mathcal{P}}) \cup \Gamma$, by its definition. Thus, no such p-cycle can be contained in $\overline{D_{\mathcal{P}}}$. Refer to Figure 8. ◀

▶ **Lemma 12.** *Suppose bisector $J(s_\alpha, s_\beta)$ appears within $R(\alpha)$ (see Figure 9). For any connected component e of $J(s_\alpha, s_\beta) \cap R(\alpha)$ that is not intersecting α, the label s_α must appear on the same side of e as α. Let $\partial R_e(\alpha)$ denote the portion of $\partial R(\alpha)$* cut out *by such a component e, at opposite side from α. Then $\partial R_e(\alpha) \subseteq D(s_\beta, s_\alpha)$.*

By Lemma 12, any components of $J(s_\alpha, s_\beta) \cap R(\alpha)$ must appear sequentially along $\partial R(\alpha)$. Note that $\partial R_e(\alpha)$ may as well contain Γ-arcs.

Proof. Suppose for the sake of contradiction that there is such a component $e \subseteq J(s_\alpha, s_\beta) \cap R(\alpha)$ with the label s_α appearing at opposite side of e as α (see Figure 10). Then e and $\partial R(\alpha)$ form an s_α-cycle C within $\overline{D_{\mathcal{P}}}$, contradicting Lemma 11. Suppose now that $\partial R_e(\alpha)$ lies only partially in $D(s_\beta, s_\alpha)$. Then $J(s_\beta, s_\alpha)$ would have to re-enter $R(\alpha)$ at $\partial R_e(\alpha)$, resulting in another component of $J(s_\beta, s_\alpha) \cap R(\alpha)$ with an invalid labeling. ◀

The following lemma extends Lemma 12 when a component e of $J(s_\alpha, s_\beta) \cap R(\alpha)$ intersects arc α. If $J(s_\alpha, s_\beta)$ intersects α, then there is also a component $\tilde{\beta}$ of $J(s, s_\beta) \cap R(\alpha)$ intersecting α at the same point as e. If $\tilde{\beta}$ has only one endpoint on α, let $\partial R_e(\alpha)$ denote the portion of $\partial R(\alpha)$ that is *cut out* by e, at the side of its s_β-label (see Figure 11(a)). If both endpoints of $\tilde{\beta}$ are on α then there are two components of $J(s_\alpha, s_\beta) \cap R(\alpha)$ incident to α (see Figure 11(b)); let $\partial R_e(\alpha)$ denote the portion of $\partial R(\alpha)$ between these two components.

▶ **Lemma 13.** *Let e be a component of $J(s_\alpha, s_\beta) \cap R(\alpha)$. Then $\partial R_e(\alpha) \subseteq D(s_\beta, s_\alpha)$.*

Using the basic property of Lemma 12 and its extension, we show that if there is any non-empty component of $J(s_\alpha, s_\beta) \cap R(\alpha)$, then $J(s, s_\beta)$ must also intersect $D_{\mathcal{P}}$, i.e., there exists a non-empty component of $J(s, s_\beta) \cap D_{\mathcal{P}}$ that is *missing from* $\mathcal{P}$. Using this property and Theorem 18 of the next section, we obtain the following theorem (see Section 5).

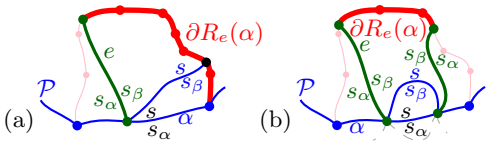

Figure 11 Illustrations for Lemma 13. The bold red parts $\partial R_e(\alpha)$ belong to $D(s_\beta, s_\alpha)$.

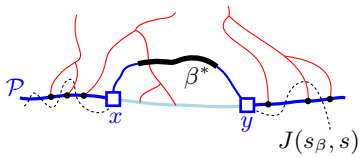

Figure 12 $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$, core arc β^* is bold, black. Endpoints of β are x, y.

▶ **Theorem 14.** *Given a boundary curve $\mathcal{P}$ of $\mathcal{S}' \subseteq \mathcal{S}$, $\mathcal{V}_l(\mathcal{P})$ (if it exists) is unique.*

The complexity of $\mathcal{V}_l(\mathcal{P})$ is $O(|\mathcal{P}|)$, where $|\mathcal{P}|$ denotes the number of boundary arcs in $\mathcal{P}$, as it is a planar graph with exactly one face per boundary arc and vertices of degree 3 (or 1).

4 Insertion in a Voronoi-like diagram

Consider a boundary curve $\mathcal{P}$ for $\mathcal{S}' \subset \mathcal{S}$ and its Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$. Let β^* be an arc in $\mathcal{S} \setminus \mathcal{S}'$, thus, β^* is contained in the closure of the domain $\overline{D_\mathcal{P}}$.

We define arc $\beta \supseteq \beta^*$ as the connected component of $J(s, s_\beta) \cap \overline{D_\mathcal{P}}$ that contains β^* (see Figure 12). We also define an insertion operation $\oplus$, which inserts arc β in $\mathcal{P}$ deriving a new boundary curve $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$, and also inserts $R(\beta)$ in $\mathcal{V}_l(\mathcal{P})$ deriving the Voronoi-like diagram $\mathcal{V}_l(\mathcal{P}_\beta) = \mathcal{V}_l(\mathcal{P}) \oplus \beta$. $\mathcal{P}_\beta$ is the boundary curve obtained by deleting the portion of $\mathcal{P}$ between the endpoints of β, which lies in $D(s_\beta, s)$, and substituting it with β.

Figure 13 enumerates the possible cases of inserting arc β in $\mathcal{P}$ and is summarized in the following observation.

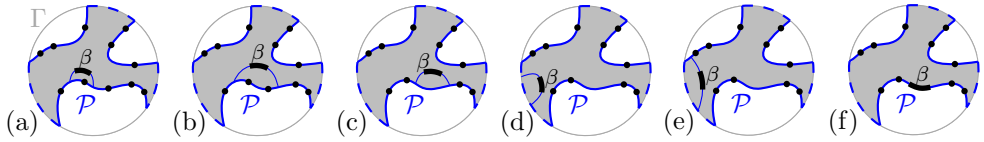

Figure 13 Insertion cases for an arc β.

▶ **Observation 15.** *Possible cases of inserting arc β in $\mathcal{P}$ (see Figure 13). $D_{\mathcal{P}_\beta} \subseteq D_\mathcal{P}$.*

(a) *β straddles the endpoint of two consecutive boundary arcs; no arcs in $\mathcal{P}$ are deleted.*
(b) *Auxiliary arcs in $\mathcal{P}$ are deleted by β; their regions are also deleted from $\mathcal{V}_l(\mathcal{P}_\beta)$.*
(c) *An arc $\alpha \in \mathcal{P}$ is split into two arcs by β; $R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$ will also be split.*
(d) *A Γ-arc is split in two by β; $\mathcal{V}_l(\mathcal{P}_\beta)$ may switch from being a tree to being a forest.*
(e) *A Γ-arc is deleted or shrunk by inserting β. $\mathcal{V}_l(\mathcal{P}_\beta)$ may become a tree.*
(f) *$\mathcal{P}$ already contains a boundary arc $\bar{\beta} \supseteq \beta^*$; then $\beta = \bar{\beta}$ and $\mathcal{P}_\beta = \mathcal{P}$.*

Note that $\mathcal{P}_\beta$ may contain fewer, the same number, or even one extra auxiliary arc compared to $\mathcal{P}$.

▶ **Lemma 16.** *The curve $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$ is a boundary curve for $\mathcal{S}' \cup \{\beta^*\}$.*

Proof. Since $\mathcal{P}$ is a (closed) s-monotone path in $\mathcal{J}_{s,\mathcal{S}'} \cup \Gamma$, $\mathcal{P}_\beta$ is also such a path in $\mathcal{J}_{s,\mathcal{S}' \cup \{s_\beta\}} \cup \Gamma$, by construction. No original arc in $\mathcal{P}$ can be deleted by the insertion of β, because every core arc in $\mathcal{S}$ appears on the envelope $\text{env}(\mathcal{J}_{s,\mathcal{S}} \cup \Gamma)$; thus, such an arc cannot be *cut out* by the insertion of β on $\mathcal{P}$. Hence, $\mathcal{P}_\beta$ contains all arcs in $\mathcal{S}' \cup \{\beta^*\}$. ◀

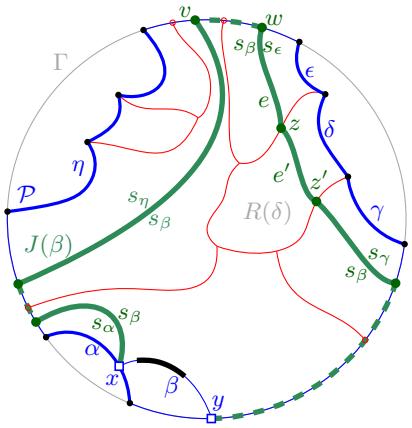

Figure 14 The merge curve $J(\beta)$ (thick, green) on $\mathcal{V}_l(\mathcal{P})$ (thin, red).

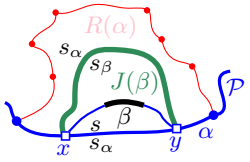

Figure 15 If β splits α, $J(\beta) \subset R(\alpha)$ would yield a forbidden s_α-inverse cycle.

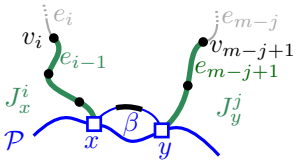

Figure 16 J^i_x and J^j_y in Section 4.1.

Given $\mathcal{V}_l(\mathcal{P})$ and arc β, where $\beta^* \in \mathcal{S} \setminus \mathcal{S}'$, we define a *merge curve* $J(\beta)$, within $\mathcal{V}_l(\mathcal{P})$, which delimits the boundary of $R(\beta)$ in $\mathcal{V}_l(\mathcal{P}_\beta)$. We define $J(\beta)$ incrementally, starting at an endpoint of β. Let x and y denote the endpoints of β, where x, β, y are in counterclockwise order around $\mathcal{P}_\beta$; refer to Figure 14.

▶ **Definition 17.** Given $\mathcal{V}_l(\mathcal{P})$ and arc $\beta \subset J(s, s_\beta)$, the *merge curve* $J(\beta)$ is a path $(v_1, \ldots, v_m)$ in the arrangement of s_β-related bisectors, $\mathcal{J}_{s_\beta, S_\mathcal{P}} \cup \Gamma$, connecting the endpoints of β, $v_1 = x$ and $v_m = y$. Each edge $e_i = (v_i, v_{i+1})$ is an arc of a bisector $J(s_\beta, \cdot)$, called an *ordinary* edge, or an arc on Γ. For $i = 1$: if $x \in J(s_\beta, s_\alpha)$, then $e_1 \subseteq J(s_\beta, s_\alpha)$; if $x \in \Gamma$, then $e_1 \subseteq \Gamma$. Given v_i, vertex v_{i+1} and edge e_{i+1} are defined as follows (see Figure 14). Wlog we assume a clockwise ordering of $J(\beta)$.

1. If $e_i \subseteq J(s_\beta, s_\alpha)$, let v_{i+1} be the other endpoint of the component $J(s_\beta, s_\alpha) \cap R(\alpha)$ incident to v_i. If $v_{i+1} \in J(s_\beta, \cdot) \cap J(s_\beta, s_\alpha)$, then $e_{i+1} \subseteq J(s_\beta, \cdot)$. If $v_{i+1} \in \Gamma$, then $e_{i+1} \subseteq \Gamma$. (In Figure 14, see $e_i = e', v_i = z, v_{i+1} = z'$.)
2. If $e_i \subseteq \Gamma$, let g be the Γ-arc incident to v_i. Let $e_{i+1} \subseteq J(s_\beta, s_\gamma)$, where $R(\gamma)$ is the first region, incident to g clockwise from v_i, such that $J(s_\beta, s_\gamma)$ intersects $g \cap \overline{R(\gamma)}$; let v_{i+1} be this intersection point. (In Figure 14, see $v_i = v$ and $v_{i+1} = w$.)

A vertex v along $J(\beta)$, is called *valid* if v is a vertex in the arrangement $\mathcal{A}(\mathcal{J}_{s_\beta, S_\mathcal{P}} \cup \Gamma)$ or v is an endpoint of β. The following theorem shows that $J(\beta)$ is well defined, given $\mathcal{V}_l(\mathcal{P})$, and that it forms an s_β-monotone path. We defer its proof to the end of this section.

▶ **Theorem 18.** *$J(\beta)$ is a unique s_β-monotone path in the arrangement of s_β-related bisectors $\mathcal{J}_{s_\beta, S_\mathcal{P}} \cup \Gamma$ connecting the endpoints of β. $J(\beta)$ can contain at most one ordinary edge per region of $\mathcal{V}_l(\mathcal{P})$, with the exception of e_1 and e_{m-1}, when v_1 and v_m are incident to the same face in $\mathcal{V}_l(\mathcal{P})$. $J(\beta)$ cannot intersect the interior of arc β.*

We define $R(\beta)$ as the area enclosed by $\beta \cup J(\beta)$. Let $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ be the subdivision of $D_{\mathcal{P}_\beta}$ obtained by inserting $J(\beta)$ in $\mathcal{V}_l(\mathcal{P})$ and deleting any portion of $\mathcal{V}_l(\mathcal{P})$ enclosed by $J(\beta)$, i.e., $\mathcal{V}_l(\mathcal{P}) \oplus \beta = \big((\mathcal{V}_l(\mathcal{P}) \setminus R(\beta)) \cup J(\beta)\big) \cap D_{\mathcal{P}_\beta}$. We prove that $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ is a Voronoi-like diagram. To this goal we need an additional property of $J(\beta)$.

▶ **Lemma 19.** *If the insertion of β splits an arc $\alpha \in \mathcal{P}$ (Observation 15(c)), then $J(\beta)$ also splits $R(\alpha)$ and $J(\beta) \not\subseteq R(\alpha)$. In no other case can $J(\beta)$ split a region $R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$.*

Proof. Suppose for the sake of contradiction that β splits arc α and $J(\beta) \subset R(\alpha)$, as shown in Figure 15. Then $J(\beta) = J(s_\alpha, s_\beta) \cap R(\alpha)$ and the bisector $J(s_\alpha, s_\beta)$ together with the arc α form a forbidden s_α-inverse cycle, deriving a contradiction to Lemma 2. Thus, $J(\beta)$ must intersect $\partial R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$ and therefore $J(\beta) \not\subseteq R(\alpha)$. By Theorem 18, $J(\beta)$ can only enter some other region at most once. Thus, $J(\beta)$ cannot split any other region. ◀

▶ **Theorem 20.** *$\mathcal{V}_l(\mathcal{P}) \oplus \beta$ is a Voronoi-like diagram for $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$, denoted $\mathcal{V}_l(\mathcal{P}_\beta)$.*

Proof. By Theorem 18, $R(\beta)$ fulfills the properties of a Voronoi-like region. Moreover, the updated boundary of any other region $R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$, which is truncated by $J(\beta)$, remains an s_α-monotone path. By Lemma 19, $J(\beta)$ cannot split a region $R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$, and thus, it cannot create a face that is not incident to α. Therefore, $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ fulfills all properties of Definition 8. ◀

The tracing of $J(\beta)$ within $\mathcal{V}_l(\mathcal{P})$, given the endpoints of β, can be done similarly to any ordinary Voronoi diagram, see e.g., [11] [2, Ch. 7.5.3] for AVDs, or [9, Ch. 7.4] [18, Ch. 5.5.2.1] for concrete diagrams. For a Voronoi-like diagram this can be established due to the basic property of Lemmas 12 and 13.

Special care is required in cases (c), (d), and (e) of Observation 15, in order to identify the first edge of $J(\beta)$; in these cases, β may not overlap with any feature of $\mathcal{V}_l(\mathcal{P})$, thus, a starting point for tracing $J(\beta)$ is not readily available. In case (c), we trace a portion of $\partial R(\alpha)$, which does not get deleted afterwards, thus it adds to the time complexity of the operation $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ (see Lemma 21). In cases (d) and (e), we show that if no feature of $\mathcal{V}_l(\mathcal{P})$ overlaps β, then either there is a leaf of $\mathcal{V}_l(\mathcal{P})$ in the neighboring Γ-arc or $J(\beta) \subseteq \overline{R(\alpha)}$. In either case a starting point for $J(\beta)$ can be identified in $O(1)$ time. Notice, if $J(\beta) \subseteq \overline{R(\alpha)}$, then it consists of a single bisector $J(s_\beta, s_\alpha)$ (and one or two Γ-arcs).

The following lemma gives the time complexity to compute $J(\beta)$ and update $\mathcal{V}_l(\mathcal{P}_\beta)$. The statement of the lemma is an adaptation from [10], however, the proof contains cases that do not appear in a farthest segment Voronoi diagram. $|\cdot|$ denotes complexity.

Let $\tilde{\mathcal{P}}$ denote a finer version of $\mathcal{P}$, where a Γ-arc between two consecutive boundary arcs in $\mathcal{P}$ is partitioned into smaller Γ-arcs as defined by the incident faces of $\mathcal{V}_l(\mathcal{P})$. Since $|\mathcal{V}_l(\mathcal{P})|$ is $O(|\mathcal{P}|)$, $|\tilde{\mathcal{P}}|$ is also $O(|\mathcal{P}|)$.

▶ **Lemma 21.** *Let α and γ be the first original arcs on $\mathcal{P}_\beta$ occurring before and after β. Let $d(\beta)$ be the number of arcs in $\tilde{\mathcal{P}}$ between α and γ (both boundary and Γ-arcs). Given α, γ, and $\mathcal{V}_l(\mathcal{P})$, in all cases of Observation 15, except (c), the merge curve $J(\beta)$ and the diagram $\mathcal{V}_l(\mathcal{P}_\beta)$ can be computed in time $O(|R(\beta)| + d(\beta))$. In case (c), where an arc is split and a new arc ω is created by the insertion of β, the time is $O(|\partial R(\beta)| + |\partial R(\omega)| + d(\beta))$.*

4.1 Proving Theorem 18

We first establish that $J(\beta)$ cannot intersect arc β, other than its endpoints, using the following Lemma.

▶ **Lemma 22.** *Given $\mathcal{V}_l(\mathcal{P})$, for any arc $\alpha \in \mathcal{P}$, $R(\alpha) \subseteq D(s, s_\alpha)$.*

Proof. The contrary would yield an s_α-inverse cycle defined by $J(s, s_\alpha)$ and $\partial R(\alpha)$. ◀

Lemma 22 implies that bisector $J(s_\beta, s_\alpha)$ cannot intersect $J(s, s_\beta)$ within region $R(\alpha)$. Thus $J(\beta)$ cannot intersect arc β in its interior. The following lemma is used in several proofs.

▶ **Lemma 23.** *$D(s, \cdot) \cap D_{\mathcal{P}}$ is always connected. Thus, any components of $J(s, \cdot) \cap D_{\mathcal{P}}$ must appear sequentially along $\mathcal{P}$.*

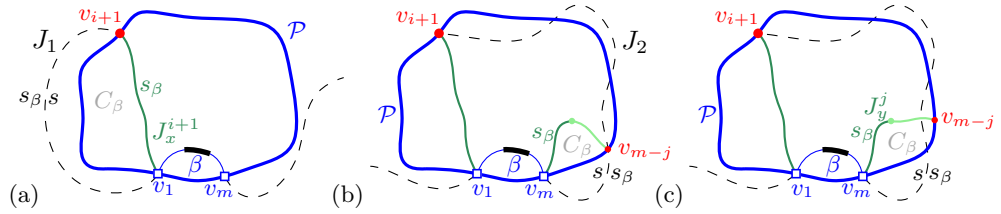

Figure 17 The assumption that edge $e_i = (v_i, v_{i+1})$ of the merge curve J_x^i hits a boundary arc of $\mathcal{P}$ as in Lemma 24.

Proof. If we assume the contrary we obtain an s-inverse cycle defined by $J(s,\cdot)$ and $\mathcal{P}$. ◀

To prove Theorem 18 we use a bi-directional induction on the vertices of $J(\beta)$. Let $J_x^i = (v_1, v_2, \ldots, v_i), 1 \le i < m$, be the subpath of $J(\beta)$ starting at $v_1 = x$ up to vertex v_i, including a small neighborhood of e_i incident to v_i, see Figure 16. Note that vertex v_i uniquely determines e_i, however, its other endpoint is not yet specified. Similarly, let $J_y^j = (v_m, v_{m-1}, \ldots, v_{m-j+1}), 1 \le j < m$, denote the subpath of $J(\beta)$, starting at v_m up to vertex v_{m-j+1}, including a small neighborhood of edge e_{m-j}. Recall that we refer to the edges of $J(\beta)$ that are not Γ-arcs as *ordinary*. For any ordinary edge $e_\ell \in J(\beta)$, let α_ℓ denote the boundary arc that *induces* e_ℓ, i.e., $e_\ell \subseteq J(s_{\alpha_\ell}, s_\beta) \cap R(\alpha_\ell)$.

Induction hypothesis: Suppose J_x^i and J_y^j, $i, j \ge 1$, are disjoint s_β-monotone paths. Suppose further that each ordinary edge of J_x^i and of J_y^j passes through a distinct region of $\mathcal{V}_l(\mathcal{P})$: α_ℓ is distinct for ℓ, $1 \le \ell \le i$ and $m - j \le \ell < m$, except possibly $\alpha_i = \alpha_{m-j}$ and $\alpha_1 = \alpha_{m-1}$.

Induction step: Assuming that $i + j < m$, we prove that at least one of J_x^i or J_y^j can respectively grow to J_x^{i+1} or J_y^{j+1} at a valid vertex (Lemmas 24, 25), and it enters a new region of $\mathcal{V}_l(\mathcal{P})$ that has not been visited so far (Lemma 27). A finish condition when $i + j = m$ is given in Lemma 26. The base case for $i = j = 1$ is trivially true.

Suppose that $e_i \subseteq J(s_{\alpha_i}, s_\beta)$ and $v_i \in \partial R(\alpha_i)$. To show that v_{i+1} is a valid vertex it is enough to show that (1) v_{i+1} can not be on α_i, and (2) if v_i is on a Γ-arc then v_{i+1} can be determined on the same Γ-arc. However, we cannot easily derive these conclusions directly. Instead we show that if v_{i+1} is not valid then v_{m-j} will have to be valid.

In the following lemmas we assume that the induction hypothesis holds.

▶ **Lemma 24.** *Suppose $e_i \subseteq J(s_{\alpha_i}, s_\beta)$ but $v_{i+1} \in \alpha_i$, i.e., it is not a valid vertex because e_i hits α_i. Then vertex v_{m-j} must be a valid vertex in $\mathcal{A}(\mathcal{J}_{s_\beta, S_\mathcal{P}})$, and v_{m-j} can not be on $\mathcal{P}$.*

▶ **Lemma 25.** *Suppose vertex v_i is on a Γ-arc g but v_{i+1} cannot be determined because no bisector $J(s_\beta, s_\gamma)$ intersects $\overline{R(\gamma)} \cap g$, clockwise from v_i. Then vertex v_{m-j} must be a valid vertex in $\mathcal{A}(\mathcal{J}_{s_\beta, S_\mathcal{P}})$ and v_{m-j} can not be on $\mathcal{P}$.*

Proof of Lemma 24. Suppose vertex v_{i+1} of e_i lies on α_i as shown in Figure 17(a). Vertex v_{i+1} is the intersection point of related bisectors $J(s, s_{\alpha_i})$, $J(s_\beta, s_{\alpha_i})$ and thus also of $J(s, s_\beta)$. Observe that arc β partitions $J(s, s_\beta)$ in two parts: J_1 incident to v_1 and J_2 incident to v_m. We claim that v_{i+1} lies on J_2. Suppose otherwise, then J_x^{i+1} and J_1 would form a forbidden s_β-inverse cycle, see the dashed black and the green solid curve in Figure 17(a). By Lemma 23 the components of $J_2 \cap D_\mathcal{P}$ appear on $\mathcal{P}$ clockwise after v_{i+1} and before v_m, as shown in Figure 17(b) illustrating $J(s, s_\beta)$ as a black dashed curve.

Now consider J_y^j. We show that v_{m-j} cannot be on $\mathcal{P}$. First observe that v_{m-j} can not lie on $\mathcal{P}$, clockwise after v_m and before v_1, since J_y^{j+1} cannot cross β. We prove that v_{m-j}

cannot lie on $\mathcal{P}$ clockwise after v_1 and before v_{i+1}. To see that, note that edge e_{m-j} cannot cross any non-Γ edge of J_x^{i+1}, because α_{m-j} is distinct from all $\alpha_\ell, \ell \leq i$ (by the induction hypothesis). In addition, by the definition of a Γ-arc, v_{m-j} cannot lie on any Γ-arc of J_x^i. Finally, we show that v_{m-j} cannot lie on $\mathcal{P}$ clockwise after v_{i+1} and before v_m. If v_{m-j} lay on the boundary arc α_{m-j} then we would have $v_{m-j} \in J(s, s_\beta)$. This would define an s_β-inverse cycle C_β, formed by J_y^{j+1} and $J(s_\beta, s)$, see Figure 17(b). If v_{m-j} lay on a Γ-arc then there would also be a forbidden s_β-inverse cycle formed by J_y^{j+1} and $J(s, s_\beta)$ because in order to reach Γ edge e_i must cross $J(s, s_\beta)$. See the dashed black and the green curve in Figure 17(c). Thus $v_{m-j} \notin \mathcal{P}$.

Since $v_{m-j} \in \partial R(\alpha_{i+1})$ but $v_{m-j} \notin \mathcal{P}$, it must be a vertex of $\mathcal{A}(\mathcal{J}_{s_\beta, S_\mathcal{P}})$. ◀

Lemma 26 provides a finish condition for the induction. When it is met, $J(\beta) = J_x^i \cup J_y^j$, i.e., a concatenation of J_x^i and J_y^j.

▶ **Lemma 26.** *Suppose $i + j > 2$ and either (1) or (2) holds: (1) $\alpha_i = \alpha_{m-j}$, i.e., v_i and v_{m-j+1} are incident to a common region $R(\alpha_i)$ and $e_i, e_{m-j} \subseteq J(s_\beta, s_{\alpha_i})$; or (2) v_i and v_{m-j+1} are on a common Γ-arc g of $\mathcal{P}$ and $e_i, e_{m-j} \subseteq \Gamma$. Then $v_{i+1} = v_{m-j+1}$, $v_{m-j} = v_i$, and $m = i + j$.*

▶ **Lemma 27.** *Suppose vertex v_{i+1} is valid and $e_{i+1} \subseteq J(s_\beta, s_{a_{i+1}})$. Then $R(\alpha_{i+1})$ has not been visited by J_x^i nor J_y^j, i.e., $\alpha_{i+1} \neq \alpha_\ell$ for $\ell \leq i$ and for $m - j < \ell$.*

By Lemma 27, J_x^{i+1} and J_y^{j+1} always enter a new region of $\mathcal{V}_l(\mathcal{P})$ that has not been visited yet; thus, conditions (1) or (2) of Lemma 26 must be fulfilled at some point of the induction. Hence, the proof of Theorem 18 is complete. Completing the induction establishes also that the conditions of Lemmas 24 and 25 can never be met, thus, no vertex of $J(\beta)$ can be on a boundary arc of $\mathcal{P}$, except its endpoints.

5 $\mathcal{V}_l(\mathcal{P})$ is unique

In this section we establish that $\mathcal{V}_l(\mathcal{P})$ is unique. To this goal we prove the following lemma and use it to prove Theorem 14.

▶ **Lemma 28.** *Suppose there is a non-empty component e of $J(s_\alpha, \cdot)$ intersecting $R(\alpha)$ in $\mathcal{V}_l(\mathcal{P})$. Then $J(s, \cdot)$ must also intersect $D_\mathcal{P}$. Further, there exists a component of $J(s, \cdot) \cap D_\mathcal{P}$, denoted as β, such that the merge curve $J(\beta)$ in $\mathcal{V}_l(\mathcal{P})$ contains e.*

Proof sketch of Theorem 14. Suppose that for a given boundary curve $\mathcal{P}$ there exist two different Voronoi-like diagrams $\mathcal{V}_l^1 \neq \mathcal{V}_l^2$. Then there must be an edge $e^1 \subseteq J(s_\beta, s_\alpha)$ of $\mathcal{V}_l^1$, such that e^1 intersects region $R^2(\alpha)$ of $\mathcal{V}_l^2$. Let edge $e \subseteq J(s_\beta, s_\alpha)$ be the component of $R^2(\alpha) \cap J(s_\beta, s_\alpha)$ overlapping with e^1. Lemma 28 yields a non-empty component β_0 of $J(s, s_\beta) \cap D_\mathcal{P}$ such that $J(\beta_0)$ on $\mathcal{V}_l^2$ contains edge e. Since $J(\beta_0)$ and $\partial R^1(\beta)$ have an overlapping component $e \cap e^1$, and they bound the regions of two different arcs $\beta_0 \neq \beta$ of site s_β, they form an s_β-cycle C. But C is contained in $D_\mathcal{P}$, deriving a contradiction to Lemma 11. ◀

6 A randomized incremental algorithm

Consider a random permutation of the set of arcs $\mathcal{S}$, $o = (\alpha_1, \ldots, \alpha_h)$. For $1 \leq i \leq h$ define $\mathcal{S}_i = \{\alpha_1, \ldots, \alpha_i\} \subseteq \mathcal{S}$ to be the subset of the first i arcs in o. Given $\mathcal{S}_i$, let $\mathcal{P}_i$ denote a boundary curve for $\mathcal{S}_i$, which induces a domain $D_i = D_{\mathcal{P}_i}$.

The randomized algorithm is inspired by the randomized, two-phase, approach of Chew [7] for the Voronoi diagram of points in convex position; however, it constructs Voronoi-like diagrams of boundary curves $\mathcal{P}_i$ within a series of shrinking domains $D_i \supseteq D_{i+1}$. The boundary curves are obtained by the insertion operation, starting with $J(s, s_{\alpha_1})$, thus, they always admit a Voronoi-like diagram. In phase 1, the arcs in $\mathcal{S}$ get deleted one by one in reverse order of o, while recording the neighbors of each deleted arc at the time of its deletion. Let $\mathcal{P}_1 = \partial(D(s, s_{\alpha_1}) \cap D_\Gamma)$ and $D_1 = D(s, s_{\alpha_1}) \cap D_\Gamma$. Let $R(\alpha_1) = D_1$. $\mathcal{V}_l(\mathcal{P}_1) = \emptyset$ is the Voronoi-like diagram for $\mathcal{P}_1$. In phase 2, we start with $\mathcal{V}_l(\mathcal{P}_1)$ and incrementally compute $\mathcal{V}_l(\mathcal{P}_{i+1})$, $i = 1, \ldots, h-1$, by inserting arc α_{i+1} in $\mathcal{V}_l(\mathcal{P}_i)$, where $\mathcal{P}_{i+1} = \mathcal{P}_i \oplus \alpha_{i+1}$ and $\mathcal{V}_l(\mathcal{P}_{i+1}) = \mathcal{V}_l(\mathcal{P}_i) \oplus \alpha_{i+1}$. At the end we obtain $\mathcal{V}_l(\mathcal{P}_h)$, where $\mathcal{P}_h = \mathcal{S}$.

We have already established that $\mathcal{V}_l(\mathcal{S}) = \mathcal{V}(\mathcal{S})$ (Corollary 10) and $\mathcal{P}_h = \mathcal{S}$, thus, the algorithm is correct. Given the analysis and the properties of Voronoi-like diagrams established in Sections 3 and 4, as well as Lemma 21, the time analysis becomes similar to the one for the farthest-segment Voronoi diagram [10].

▶ **Lemma 29.** *$\mathcal{P}_i$ contains at most $2i$ arcs; thus, the complexity of $\mathcal{V}_l(\mathcal{P}_i)$ is $O(i)$.*

Proof. At each step of phase 2, one original arc is inserted and at most one additional arc is created by a split, thus, $|\mathcal{P}_i| \leq 2i$. The complexity of $\mathcal{V}_l(\mathcal{P}_i)$ is $O(|\mathcal{P}_i|)$, thus, it is $O(i)$. ◀

▶ **Lemma 30.** *The expected number of arcs in $\tilde{\mathcal{P}}_i$ (auxiliary boundary arcs and fine Γ-arcs) that are visited while inserting α_{i+1} is $O(1)$.*

Proof. To insert arc α_{i+1} at one step of phase 2, we may trace a number of arcs in $\tilde{\mathcal{P}}_i$ that may be auxiliary arcs and/or fine Γ-arcs between the pair of consecutive original arcs that has been stored with α_{i+1} in phase 1. Since every element of $\mathcal{S}_{i+1}$ is equally likely to be α_{i+1}, each pair of consecutive original arcs in $\mathcal{P}_{i+1}$ has probability $1/i$ to be considered at step i. Let n_j be the number of arcs inbetween the jth pair of original arcs in $\tilde{\mathcal{P}}_i$, $1 \leq j \leq i$; $\sum_{j=1}^{i} n_j = |\tilde{\mathcal{P}}_i|$ which is $O(i)$. The expected number of arcs that are traced is then $\sum_{j=1}^{i} n_j/i \in O(1)$. ◀

Using the same backwards analysis as in [10], we conclude with the following theorem.

▶ **Theorem 31.** *Given an abstract Voronoi diagram $\mathcal{V}(S)$, $\mathcal{V}(S \setminus \{s\}) \cap VR(s, S)$ can be computed in expected $O(h)$ time, where h is the complexity of $\partial VR(s, S)$. Thus, $\mathcal{V}(S \setminus \{s\})$ can also be computed in expected time $O(h)$.*

7 Concluding remarks

Updating an abstract Voronoi diagram, after deletion of one site, in deterministic linear time remains an open problem. We plan to investigate the applicability of Voronoi-like diagrams in the linear-time framework of Aggarwal et al. [1] in subsequent research.

The algorithms and the results in this paper (Theorem 31) are also applicable to concrete Voronoi diagrams of line segments and planar straight-line graphs (including simple and non-simple polygons) even though they do not strictly fall under the AVD model unless segments are disjoint. For intersecting line segments, $\partial\mathrm{VR}(s, S)$ is a Davenport-Schinzel sequence of order 4 [17] but this does not affect the complexity of the algorithm, which remains linear.

Examples of concrete diagrams that fall under the AVD umbrella and thus can benefit from our approach include [6]: disjoint line segments and disjoint convex polygons of constant size in the L_p norms, or under the Hausdorff metric; point sites in any convex distance metric or the Karlsruhe metric; additively weighted points that have non-enclosing circles; power diagrams with non-enclosing circles.

References

1. A. Aggarwal, L. Guibas, J. Saxe, and P. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete & Computational Geometry*, 4:591–604, 1989.
2. F. Aurenhammer, R. Klein, and D.-T. Lee. *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, 2013. URL: `http://www.worldscientific.com/worldscibooks/10.1142/8685`.
3. C. Bohler, P. Cheilaris, R. Klein, C. H. Liu, E. Papadopoulou, and M. Zavershynskyi. On the complexity of higher order abstract Voronoi diagrams. *Computational Geometry: Theory and Applications*, 48(8):539–551, 2015. `doi:10.1016/j.comgeo.2015.04.008`.
4. C. Bohler, R. Klein, and C. Liu. Forest-like abstract Voronoi diagrams in linear time. In *Proc. 26th Canadian Conference on Computational Geometry (CCCG)*, 2014. URL: `http://www.cccg.ca/proceedings/2014/papers/paper20.pdf`.
5. C. Bohler, R. Klein, and C.-H. Liu. An Efficient Randomized Algorithm for Higher-Order Abstract Voronoi Diagrams. In *32nd International Symposium on Computational Geometry (SoCG)*, volume 51, pages 21:1–21:15, Dagstuhl, Germany, 2016. `doi:10.4230/LIPIcs.SoCG.2016.21`.
6. C. Bohler, C. H. Liu, E. Papadopoulou, and M. Zavershynskyi. A randomized divide and conquer algorithm for higher-order abstract Voronoi diagrams. *Computational Geometry: Theory and Applications*, 59(C):26–38, 2016. `doi:10.1016/j.comgeo.2016.08.004`.
7. L. P. Chew. Building Voronoi diagrams for convex polygons in linear expected time. Technical report, Dartmouth College, Hanover, USA, 1990.
8. F. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete & Computational Geometry*, 21(3):405–420, 1999.
9. M. de Berg, O. Schwarzkopf, M. van Kreveld, and M. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, 2nd edition, 2000.
10. E. Khramtcova and E. Papadopoulou. An expected linear-time algorithm for the farthest-segment Voronoi diagram. arXiv:1411.2816v3 [cs.CG], 2017. Preliminary version in *Proc. 26th Int. Symp. on Algorithms and Computation (ISAAC), LNCS* 9472, 404-414, 2015.
11. R. Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes in Computer Science*. Springer-Verlag, 1989.
12. R. Klein, E. Langetepe, and Z. Nilforoushan. Abstract Voronoi diagrams revisited. *Computational Geometry: Theory and Applications*, 42(9):885–902, 2009.
13. R. Klein and A. Lingas. Hamiltonian abstract Voronoi diagrams in linear time. In *Algorithms and Computation, 5th International Symposium, (ISAAC)*, volume 834 of *Lecture Notes in Computer Science*, pages 11–19, 1994.
14. R. Klein, K. Mehlhorn, and S. Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Computational geometry: Theory and Applications*, 3:157–184, 1993.
15. K. Mehlhorn, S. Meiser, and R. Rasch. Furthest site abstract Voronoi diagrams. *International Journal of Computational Geometry and Applications*, 11(6):583–616, 2001.
16. A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley, second edition, 2000.
17. E. Papadopoulou and M. Zavershynskyi. The higher-order Voronoi diagram of line segments. *Algorithmica*, 74(1):415–439, 2016. `doi:10.1007/s00453-014-9950-0`.
18. S. M. Preparata F.P. *Computational Geometry*. Texts and Monographs in Computer Science. Springer, New York, NY, 1985.
19. M. Sharir and P. K. Agarwal. *Davenport-Schinzel sequences and their geometric applications*. Cambridge university press, 1995.