



Deletion in Abstract Voronoi Diagrams in Expected Linear Time and Related Problems

Kolja Junginger¹ · Evanthia Papadopoulou¹

Received: 21 July 2020 / Revised: 16 May 2022 / Accepted: 23 September 2022

© The Author(s) 2023

Abstract

Updating an abstract Voronoi diagram in linear time, after deletion of one site, has been an open problem in a long time; similarly, for any concrete Voronoi diagram of generalized (non-point) sites. In this paper we present a simple, expected linear-time algorithm to update an abstract Voronoi diagram after deletion of one site. To achieve this result, we use the concept of a Voronoi-like diagram, a relaxed Voronoi structure of independent interest. 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 insertion, therefore, enabling its use in incremental constructions. The time-complexity analysis introduces a variant to backwards analysis, which is applicable to order-dependent structures. We further extend the technique to compute in expected linear time: the order- $(k + 1)$ subdivision within an order- k Voronoi region, and the farthest abstract Voronoi diagram, after the order of its regions at infinity is known.

Keywords Abstract Voronoi diagram · Linear-time algorithm · Randomized incremental construction · Backwards analysis · Site-deletion · Higher-order Voronoi diagram · Farthest Voronoi diagram

Mathematics Subject Classification 68W05 · 68U05

Editor in Charge: Kenneth Clarkson

This research was supported in part by the Swiss National Science Foundation, project 200021E_154387. A preliminary version of this paper appeared in *Proc. 34th International Symposium on Computational Geometry (SoCG) 2018*.

Kolja Junginger
junginger.kolja@gmail.com

Evanthia Papadopoulou
evanthia.papadopoulou@usi.ch

¹ Faculty of Informatics, USI Università della Svizzera italiana, Lugano, Switzerland

1 Introduction

The Voronoi diagram of a set S of n simple geometric objects, called sites, is a versatile geometric partitioning structure that reveals proximity information among the input sites. Classic variants include the *nearest-neighbor*, the *farthest-site*, and the *order- k* Voronoi diagram of the set S . Abstract Voronoi diagrams [11] offer a unifying framework to many concrete and fundamental instances. Voronoi diagrams have been well investigated and many optimal construction algorithms exist in various cases. For more information, see, e.g., the book of Aurenhammer et al. [2], and the book of Okabe et al. [17] for a wealth of applications.

For certain Voronoi diagrams with a tree structure, linear-time algorithms have been well known to exist for their construction, see e.g., [1, 7, 8, 13]. The first linear-time 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. The same technique 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 time linear to the number of Voronoi neighbors of the deleted site; (2) computing the order- $(k+1)$ subdivision within an order- k Voronoi region; (3) computing the farthest Voronoi diagram of point-sites in linear time, given their convex hull. A much simpler randomized technique for the same problems was introduced by Chew [7]. The medial axis of a simple polygon is another well-known problem that admits a linear-time construction, as shown by Chin et al. [8].

Surprisingly, no linear-time constructions have been known for any of the problems (1)–(3) for Voronoi diagrams involving non-point sites, and similarly for abstract Voronoi diagrams. Under restrictions, Klein and Lingas [13] adapted the linear-time approach of [1] to the abstract framework 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 some forest structures within a given domain [4], under similar restrictions, where no region can have multiple faces and each bisector can intersect this domain in one component.

In this paper we consider the fundamental problem of site-deletion in abstract Voronoi diagrams and provide a simple expected linear-time technique to achieve this task. We work in the framework of abstract Voronoi diagrams so that we can simultaneously address all the concrete instances that fall under their umbrella. After deletion (1), we extend the randomized linear-time technique to the remaining problems: (2) computing the order- $(k+1)$ subdivision within an order- k abstract Voronoi region; and (3) computing the farthest abstract Voronoi diagram after the order of its faces at infinity is known. The latter sequence of faces can be computed in time $O(n \log n)$. To the best of our knowledge, no deterministic linear-time technique is yet known for these problems.

To achieve our goal, we define the *Voronoi-like diagram*, a relaxed Voronoi structure, which is interesting in its own right. *Voronoi-like regions* are supersets of real Voronoi regions, and their boundaries correspond to simple *monotone paths* in the arrangement of the underlying bisector system (see Definition 3.1). We prove the correctness and




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

uniqueness of this structure, and use it to derive a simple randomized incremental algorithm to address the above problems in linear expected time.

An earlier attempt towards a linear-time construction for the farthest-segment Voronoi diagram appeared in [10] following a different geometric formulation, which does not extend to the abstract setting. A preliminary version of the present paper, regarding site deletion in abstract Voronoi diagrams, appeared in [9]. In three dimensions, site-deletion in Delaunay triangulations of point-sites, as inspired by the randomized approach of Chew [7], has been considered in [6].

Abstract Voronoi diagrams (AVDs). These diagrams were introduced by Klein [11]. Instead of sites and distance measures, they are defined in terms of bisector 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 Fig. 1. The *Voronoi region of p* is

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

The (*nearest-neighbor*) *Voronoi diagram of S* is

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

Following the traditional model of AVDs (see, e.g., [3, 4, 11]) the bisector system is assumed to satisfy the following axioms, for every subset $S' \subseteq S$:

- (A1) Each Voronoi region $VR(p, S')$ is non-empty and path-connected.
- (A2) Each point in the plane belongs to the closure of a Voronoi region $VR(p, S')$.
- (A3) Each bisector $J(p, q)$ is an unbounded curve, which after stereographic projection to the sphere can be completed to a closed 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]).

The abstract Voronoi diagram $\mathcal{V}(S)$ is a plane graph of structural complexity $O(n)$ whose 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 need to compute $\mathcal{V}(S \setminus \{s\})$ within $VR(s, S)$. This diagram is a tree, if $VR(s, S)$ is bounded, and a forest otherwise.

However, its regions can be disconnected, i.e., one region may consist of multiple faces. The site-occurrences along $\partial \text{VR}(s, S)$ form a Davenport–Schinzel sequence of order 2. Disconnected regions introduce severe complications which differentiate the problem from its counterpart on point-sites. For example, let $S' \subset S \setminus \{s\}$; the diagram $\mathcal{V}(S') \cap \text{VR}(s, S' \cup \{s\})$ may contain faces that do not even appear in $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$, and conversely, an arbitrary sub-sequence of arcs on $\partial \text{VR}(s, S)$ need not be related to any Voronoi diagram of sites in S . At a first sight, a linear-time algorithm may seem infeasible.

Our results. In this paper we formalize the concept of a *Voronoi-like diagram*, a relaxed Voronoi structure defined as an acyclic graph (a tree or forest) in the arrangement of the underlying bisector system, and prove that it is well defined and unique. This structure provides a tool to deal with disconnected Voronoi regions, and thus, address the site-deletion problem efficiently. We envision that it will be useful in other cases of Voronoi diagrams with disconnected regions as well.

Given a Voronoi-like diagram, we define an *insertion operation* and prove its correctness. This makes a simple randomized incremental construction possible. The time analysis of the randomized algorithm is non-standard because the intermediate Voronoi-like structures are order-dependent. We give a technique, which offers a simple variant to backwards analysis that can be applied to order-dependent structures. We partition the set of permutations of length i into manageable groups of i permutations each, and show that the time complexity of step i in each group is $O(i)$. We can then conclude that step i is performed in expected $O(1)$ time.

In this paper we focus on site-deletion, and compute $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ in expected time linear in the number of Voronoi neighbors of the deleted site. We also extend the approach to address the aforementioned related problems for the order- k and the farthest abstract Voronoi diagram, problems (2) and (3), respectively.

Examples of concrete diagrams that fall under the AVD umbrella, and thus, can benefit from our approach include: 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.

This paper is organized as follows. Section 2 provides background on abstract Voronoi diagrams. Section 3 formulates the Voronoi-like diagram, which is implied by a subset of $\partial \text{VR}(s, S)$, given a fixed site $s \in S$. Section 4 defines an insertion operation on a Voronoi-like diagram and proves its correctness. Section 5 proves the uniqueness of the Voronoi-like diagram of a boundary curve. Section 6 outlines the simple randomized incremental construction and proves its time complexity. To this goal, Sect. 6.1 gives a variant of backwards analysis that is applicable to order-dependent structures. To follow the algorithm in Sect. 6 only the basic definitions in Sect. 3 are needed; the correctness and uniqueness proofs of the previous sections are not necessary to follow the algorithm, and thus, they can be skipped. Sections 7 and 8 extend the technique further to the order- k and farthest abstract Voronoi diagram respectively. Section 9 gives concluding remarks.




Fig. 2 The Voronoi diagram of three sites, if related bisectors (dashed lines) intersect twice; $\text{VR}(p, \{p, q, r\})$ is shown shaded

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 $\mathcal{J} = \{J(p, q) : p \neq q \in S\}$. \mathcal{J} fulfills axioms (A1)–(A4), as given in Sect. 1, for every $S' \subseteq S$.

Bisectors in \mathcal{J} that have a site p in common are called *p-related* or simply *related*. Any two related bisectors can intersect at most twice [11, Lem. 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), which are the Voronoi vertices of the diagram $\mathcal{V}(\{p, q, r\})$. The Voronoi diagram of three sites $\mathcal{V}(\{p, q, r\})$ may have at most two Voronoi vertices, see Fig. 2. The set of all p -related bisectors that involve sites in any $S' \subseteq S$ is denoted $\mathcal{J}_{p, S'} = \{J(p, q) : q \in S', q \neq p\}$.

Let $\text{VR}(s, S)$ be the Voronoi region of a site $s \in S$. Although $\text{VR}(s, S)$ is simply connected, the sites in $S \setminus \{s\}$ appearing along the boundary $\partial \text{VR}(s, S)$ may repeat, forming a Davenport–Schinzel sequence of order 2. This is because s -related bisectors can intersect at most twice, and thus, [21, Thm. 5.7] applies. This is a fundamental difference from the classic case of point-sites in the Euclidean plane, where bisectors are straight-lines, therefore, they intersect at most once, and no site repetition can occur along the boundary of a Voronoi region.

Suppose we delete the site $s \in S$ from $\mathcal{V}(S)$. To update the Voronoi diagram after the deletion of s , we need to compute $\mathcal{V}(S \setminus \{s\})$ within the Voronoi region $\text{VR}(s, S)$, i.e., compute $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$. We first characterize the structure of this diagram in the following lemma. An alternative proof can also be derived from the order- k counterpart [5], which appeared after the preliminary version of this paper [9].

Lemma 2.1 $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ is a forest having exactly one face for each Voronoi edge of $\partial \text{VR}(s, S)$. Its leaves are the Voronoi vertices of $\partial \text{VR}(s, S)$, and points at infinity, if $\text{VR}(s, S)$ is unbounded (see Fig. 3). If $\text{VR}(s, S)$ is bounded then $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ is a tree.

Proof Every face in $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ must touch the boundary $\partial \text{VR}(s, S)$ because Voronoi regions are non-empty and connected; this implies that the diagram is a forest. Every Voronoi edge $e \subseteq J(s, p)$ on $\partial \text{VR}(s, S)$ must be entirely in $\text{VR}(p, S \setminus \{s\})$. Thus, no leaf can lie in the interior of a Voronoi edge of $\partial \text{VR}(s, S)$. On the other hand, each Voronoi vertex of $\partial \text{VR}(s, S)$ must be a leaf of the diagram as its incident edges are induced by different sites.

Now we show that no two edges of $\partial \text{VR}(s, S)$ can be incident to the same face of $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$. Consider two edges on $\partial \text{VR}(s, S)$ induced by the same




Fig. 3 $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ in red; $\partial \text{VR}(s, S)$ is shown in bold black




Fig. 4 $\text{VR}(p, S \setminus \{s\}) \cap \text{VR}(s, S)$ cannot be connected because of $J(p, q)$

site $p \in S \setminus \{s\}$. Then there exists an edge between them, induced by a site $q \neq p$, such that the bisector $J(s, q)$ has exactly two intersections with $J(p, s)$ as shown in Fig. 4. The bisector $J(p, q)$ intersects with them at the same two points. Since the bisector system is admissible, and thus $\text{VR}(p, \{s, p, q\})$ is connected, $J(p, q)$ connects these endpoints through $D(p, s) \cap D(q, s)$ as shown in Fig. 4, thus, $J(p, q) \cap \text{VR}(s, \{s, p, q\})$ consists of two unbounded connected components. This implies that $D(p, q) \cap \text{VR}(s, S)$ must have two disjoint faces, each of which is incident to exactly one of the two edges of p . Thus, $\text{VR}(p, S \setminus \{s\}) \cap \text{VR}(s, S)$ cannot be connected and the two edges of p must be incident to different faces of $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$.

If $\text{VR}(s, S)$ is unbounded, two consecutive edges of $\partial \text{VR}(s, S)$ can extend to infinity, in which case there is at least one edge of $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ extending to infinity between them; thus, leaves can be points at infinity. If $\text{VR}(s, S)$ is bounded, all leaves of $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ must lie on $\partial \text{VR}(s, S)$. Since no face is incident to more than one edge of $\partial \text{VR}(s, S)$, in this case $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ cannot be disconnected, and thus is a tree. \square

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 intersects Γ exactly twice and transversally. To avoid dealing with infinity, and without any loss of generality, we restrict all computations within Γ .¹ The curve Γ can be interpreted as $J(p, s_\infty)$, for any $p \in S$, where s_∞ is an additional site at infinity. Let D_Γ denote the portion of the plane enclosed by Γ . The domain of computation is $\text{VR}(s, S) \cap D_\Gamma$ and Fig. 5 illustrates possible cases.

We first make some observations regarding an admissible bisector system, which we then use as tools in the proofs throughout this paper.

Definition 2.2 Let C_p be a cycle of p -related bisectors in the arrangement of bisectors $\mathcal{J} \cup \{\Gamma\}$, see Fig. 6. If the label p appears inside the cycle, for every edge of C_p , then

¹ The presence of Γ is conceptual and its exact position unknown; we never compute coordinates on Γ .




Fig. 5 The domain of computation $VR(s, S) \cap D_\Gamma$ (shaded)



Fig. 6 a A p -inverse cycle. b A p -cycle

C_p is called a p -cycle. If the label p appears on the outside of the cycle for every edge in C_p , then C_p is called p -inverse.

Recall that Γ can be considered a p -related bisector, for any site $p \in S$, where the label p is in the interior of Γ . Thus, a p -cycle may contain arcs of Γ , while a p -inverse cycle cannot contain any Γ arcs.

Lemma 2.3 *In an admissible bisector system there is no p -inverse cycle.*

Proof Suppose a p -inverse cycle exists in the admissible bisector system. Let C_p denote a minimal such cycle, where no p -related bisector may intersect the interior of the cycle, which is denoted by D_p . Such a minimal cycle must exist, because if a bisector $J(p, q)$ intersects D_p , then it defines another (smaller) p -inverse cycle that is contained in $C_p \cup D_p$, whose interior is not intersected by $J(p, q)$. Let $S' \subseteq S$ denote the set of sites that define the edges of C_p . Considering S' , the farthest Voronoi region of p is $FVR(p, S') = \bigcap_{q \in S' \setminus \{p\}} D(q, p)$. By its definition, D_p must be identical to one face of $FVR(p, S')$. Since farthest Voronoi regions must be unbounded [3, 16], we derive a contradiction. \square

The following *transitivity lemma* is a consequence of transitivity of dominance regions [3, Lem. 2] and the fact that bisectors $J(p, q), J(q, r), J(p, r)$ intersect at the same point(s). Let \bar{X} denote the closure of a region X .

Lemma 2.4 *Suppose $z \in \mathbb{R}^2$ 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.




Fig. 7 Illustration of $\mathcal{S} = \partial \text{VR}(s, S)$ in bold (black) and $\mathcal{V}(\mathcal{S})$ in red; $\mathcal{S} = (\alpha, \beta, \gamma, \delta, \epsilon, \zeta, \eta, \vartheta)$

3 Problem Formulation, Definitions and Properties

Consider the Voronoi region $\text{VR}(s, S)$ for a fixed site $s \in S$. Let \mathcal{S} denote the sequence of Voronoi edges on the boundary of this region within the domain D_Γ , i.e., $\mathcal{S} = \partial \text{VR}(s, S) \cap D_\Gamma$. \mathcal{S} is a cyclically ordered set of arcs, where each arc is a piece of an s -related bisector defining a Voronoi edge on the boundary of $\text{VR}(s, S)$. The arcs in \mathcal{S} are called *core arcs*. Note that a single site in $S \setminus \{s\}$ may induce several of the core arcs in \mathcal{S} . For any arc $\alpha \in \mathcal{S}$, let s_α denote the site in S such that $\alpha \subseteq J(s, s_\alpha)$.

We interpret the core arcs in \mathcal{S} as sites that induce a Voronoi diagram $\mathcal{V}(\mathcal{S})$ such that $\mathcal{V}(\mathcal{S}) = \mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S) \cap D_\Gamma$, see Fig. 7. By Lemma 2.1, each face of $\mathcal{V}(\mathcal{S})$ is incident to exactly one core arc in \mathcal{S} ; thus, it can be interpreted as the Voronoi region of its incident core arc. Then, $\mathcal{V}(\mathcal{S})$ can be viewed as the Voronoi diagram of the arcs in \mathcal{S} .

The arrangement of a bisector set $\mathcal{J}' \subseteq \mathcal{J}$ is denoted by $\mathcal{A}(\mathcal{J}')$. A path P in the arrangement $\mathcal{A}(\mathcal{J}')$ is a connected sequence of alternating edges and vertices in this arrangement. An arc α of P (denoted as $\alpha \in P$) is a maximally connected collection of consecutive edges and vertices of the arrangement along P that 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. Any two consecutive arcs in P are pieces of different bisectors.

Consider the arrangement of a set of p -related bisectors $\mathcal{J}_{p,S'}, S' \subseteq S$. Since it may consist of several connected components, we also include Γ in this arrangement to unify the various components, deriving $\mathcal{A}(\mathcal{J}_{p,S'} \cup \{\Gamma\})$.

Definition 3.1 A path in the arrangement of p -related bisectors $\mathcal{J}_{p,S'} \cup \{\Gamma\}$, $S' \subseteq S$, is called *p-monotone* (or simply *monotone*) if any two consecutive arcs α, β on this path, where $\alpha \subseteq J(p, s_\alpha)$ and $\beta \subseteq J(p, s_\beta)$, coincide (within a neighborhood of their common endpoint) with two Voronoi edges of $\partial \text{VR}(p, \{p, s_\alpha, s_\beta\})$ (see Figs. 8, 9).

The boundary of the Voronoi region $\text{VR}(p, S' \cup \{p\}) \cap D_\Gamma$, $S' \subseteq S$, is an example of such a p -monotone path, which is called the *envelope* of $\mathcal{J}_{p,S'} \cup \{\Gamma\}$. Figure 9 illustrates examples of p -monotone paths, where the envelope is shown in Fig. 9a.

Definition 3.2 Consider $S' \subseteq \mathcal{S}$ and let $S' = \{s_\alpha \in S : \alpha \in S'\}$ be the sites in $S \setminus \{s\}$ that define the arcs in S' . A *boundary curve* \mathcal{P} for S' is a closed s -monotone path in the




Fig. 8 a Arcs α, β fulfill the p -monotone path condition; they do not fulfill it in b and c



Fig. 9 p -monotone paths in $\mathcal{J}_{p, \{q, r, t\}}$. a illustrates the envelope \mathcal{E} of $\mathcal{J}_{p, \{q, r, t\}}$

arrangement of s -related bisectors $\mathcal{J}_{s, S' \cup \{\Gamma\}}$ such that all arcs in S' are contained in \mathcal{P} . The open portion of the plane enclosed by \mathcal{P} is called the *domain* of \mathcal{P} , denoted $D_{\mathcal{P}}$. Given \mathcal{P} , let $S_{\mathcal{P}} = S'$.

A set $S' \subset S$ can admit several different boundary curves, see e.g., the different p -monotone paths in Fig. 9. One such boundary curve is the boundary of $\text{VR}(s, S' \cup \{s\}) \cap D_{\Gamma}$, which is called the *envelope* of S' , $\mathcal{E} = \partial \text{VR}(s, S' \cup \{s\}) \cap D_{\Gamma}$. The full set S can have only one boundary curve, which is the boundary of $\text{VR}(s, S) \cap D_{\Gamma}$. Recall that S is ordered according to $\partial \text{VR}(s, S)$, and the same ordering applies to any subset (equiv. subsequence) $S' \subset S$. Figure 10 illustrates a boundary curve for a subset of core arcs from Fig. 7.

A boundary curve \mathcal{P} on $S' \subseteq S$ consists of pieces of s -related bisectors 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 containing a core arc of S' are called *original* and others, which contain no core arc, are called *auxiliary*. Original boundary arcs in \mathcal{P} are expanded versions of the core arcs in S' . To distinguish between an original arc α and its core sub-arc in S' , we use an $*$ to denote the latter. Figure 10 illustrates a boundary curve \mathcal{P} on $S' \subseteq S$ consisting of five original arcs, one auxiliary arc (arc β') and one Γ -arc (arc g); the core arcs are illustrated in bold and the set S is shown in Fig. 7. Let $|\mathcal{P}|$ denote the number of boundary arcs in \mathcal{P} .

We now define the *Voronoi-like diagram* of a boundary curve \mathcal{P} on $S' \subseteq S$. Recall that $S' = \{s_{\alpha} \in S \setminus \{s\} \mid \alpha \in S'\}$ is the set of sites in $S \setminus \{s\}$, which define the core arcs in S' .

Definition 3.3 Given a boundary curve \mathcal{P} on $S' \subseteq S$, the *Voronoi-like diagram* of \mathcal{P} , denoted $\mathcal{V}_l(\mathcal{P})$, is a plane graph defined on the arrangement of the bisector system $\mathcal{J}_{s, S'}$ that subdivides the domain $D_{\mathcal{P}}$ as follows (see Fig. 10):

- for each boundary arc $\alpha \in \mathcal{P} \setminus \Gamma$, there is exactly one distinct face $R(\alpha, \mathcal{P})$, whose boundary is an s_{α} -monotone path in $\mathcal{J}_{s_{\alpha}, S' \cup \Gamma}$, plus arc α ;
- the faces cover the domain $D_{\mathcal{P}}$: $\bigcup_{\alpha \in \mathcal{P} \setminus \Gamma} \overline{R(\alpha, \mathcal{P})} = \overline{D_{\mathcal{P}}}$.

If the boundary curve \mathcal{P} coincides with the envelope $\mathcal{E} = \partial \text{VR}(s, S' \cup \{s\}) \cap D_{\Gamma}$, then $\mathcal{V}_l(\mathcal{P})$ is the ordinary Voronoi diagram of S' as truncated within the domain of \mathcal{E} . That




Fig. 10 A boundary curve \mathcal{P} on $S' \subseteq S$, where the core arcs in S' are shown in bold, and its Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$ is shown in red. The gray arc g is a Γ -arc, and the blue arc β' is an auxiliary arc; the remaining arcs are original. The set of core arcs \mathcal{S} is shown in Fig. 7

is, $\mathcal{V}_l(\mathcal{P}) = \mathcal{V}_l(\mathcal{E}) = \mathcal{V}(S') \cap D_{\mathcal{E}}$ (see Lemma 3.4 and Corollary 3.5 in the sequel). For an arbitrary boundary curve \mathcal{P} , the Voronoi-like regions in $\mathcal{V}_l(\mathcal{P})$ are related to the real Voronoi regions in $\mathcal{V}(S') \cap D_{\mathcal{E}}$ as supersets (see the following lemma).

Let $\mathcal{V}(\mathcal{E}) = \mathcal{V}(S') \cap D_{\mathcal{E}}$. Any face of the Voronoi diagram $\mathcal{V}(\mathcal{E})$ incident to a boundary arc $\alpha \in \mathcal{E}$ is regarded as the Voronoi region $\text{VR}(\alpha, \mathcal{E})$. We show that $R(\alpha, \mathcal{E}) = \text{VR}(\alpha, \mathcal{E})$, thus, $\mathcal{V}(\mathcal{E}) = \mathcal{V}_l(\mathcal{E})$.

Lemma 3.4 *Let \mathcal{P} be a boundary curve on $S' \subseteq S$ and let \mathcal{E} be the envelope of S' , $\mathcal{E} = \partial \text{VR}(s, S' \cup \{s\}) \cap D_{\Gamma}$. Let $\alpha \in \mathcal{P}$ and $\tilde{\alpha} \in \mathcal{E}$ be two overlapping arcs where $\alpha, \tilde{\alpha} \subseteq J(s, s_{\alpha})$. Then, $R(\alpha, \mathcal{P}) \supseteq \text{VR}(\tilde{\alpha}, \mathcal{E})$. Further, if α and $\tilde{\alpha}$ are original, i.e., $\alpha \supseteq \tilde{\alpha} \supseteq \alpha^*$, where $\alpha^* \in S'$, then $R(\alpha, \mathcal{P}) \supseteq \text{VR}(\tilde{\alpha}, \mathcal{E}) \supseteq \text{VR}(\alpha^*, S)$.*

Proof By the definition of a boundary curve, it holds that $\alpha \supseteq \tilde{\alpha}$. By the definition of a Voronoi region, bisector $J(s_{\alpha}, \cdot)$ cannot appear in the interior of any Voronoi region in $\mathcal{V}(S') \cap D_{\mathcal{E}} = \mathcal{V}(\mathcal{E})$. Since $\alpha \supseteq \tilde{\alpha}$, by the definition of a Voronoi-like region, it follows that $R(\alpha, \mathcal{P}) \supseteq \text{VR}(\tilde{\alpha}, \mathcal{E})$. Suppose that α and $\tilde{\alpha}$ are original; since $S' \subseteq S$, by the monotonicity property of Voronoi regions, we have $\text{VR}(\tilde{\alpha}, \mathcal{E}) \supseteq \text{VR}(\alpha^*, S)$. \square

As an example, refer to the Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$ of Fig. 10 versus the Voronoi diagram $\mathcal{V}(S)$ in Fig. 7: the Voronoi-like region $R(\eta, \mathcal{P})$ is a superset of the Voronoi region $\text{VR}(\eta^*, S)$ in Fig. 7; similarly $R(\alpha, \mathcal{P}) \supseteq \text{VR}(\alpha^*, S)$.

Another implication of Lemma 3.4 is that the adjacencies of the Voronoi diagram $\mathcal{V}(\mathcal{E})$, among the original arcs of \mathcal{E} , are all preserved in $\mathcal{V}_l(\mathcal{P})$ (see Figs. 7, 10). If $\mathcal{P} = \mathcal{E}$, then $\mathcal{V}_l(\mathcal{E})$ and $\mathcal{V}(\mathcal{E})$ coincide as a direct consequence of Lemma 3.4.

Corollary 3.5 $\mathcal{V}_l(\mathcal{E}) = \mathcal{V}(S') \cap D_{\mathcal{E}} = \mathcal{V}(\mathcal{E})$ for the envelope \mathcal{E} of $S' \subseteq S$.

In the remainder of this section we give basic properties of Voronoi-like regions involving their interaction with the bisectors in \mathcal{J} , which we later use in subsequent sections to derive correctness and establish that the Voronoi-like diagram is well defined.




Fig. 11 Various cases of Lemma 3.8. The shaded region illustrates $\text{cut}(e) \subseteq D(s_\beta, s_\alpha)$

3.1 Properties of Voronoi-Like Regions

The following property establishes that a Voronoi-like region $R(\alpha, \mathcal{P})$ cannot be intersected by $J(s, s_\alpha)$.

Lemma 3.6 For any arc $\alpha \in \mathcal{P}$, $R(\alpha, \mathcal{P}) \subseteq D(s, s_\alpha)$.

Proof The contrary would yield a forbidden s_α -inverse cycle defined by a component of $J(s, s_\alpha) \cap R(\alpha, \mathcal{P})$ and the incident portion of $\partial R(\alpha, \mathcal{P})$. \square

Lemma 3.7 For a boundary curve \mathcal{P} , its domain $\overline{D_\mathcal{P}}$ may not contain a p -cycle formed by the bisectors of $\mathcal{J}_{s, S_\mathcal{P}} \cup \{\Gamma\}$ for any site $p \in S_\mathcal{P}$.

Proof Let $p \in S_\mathcal{P}$. Any original arc of p in \mathcal{P} is bounding $\text{VR}(p, S_\mathcal{P} \cup \{s\})$, thus, it must have a portion within the interior of $\text{VR}(p, S_\mathcal{P})$ in $\mathcal{V}(S_\mathcal{P})$. Hence, $\text{VR}(p, S_\mathcal{P})$ must have some non-empty portion outside the closure of $D_\mathcal{P}$. However, $\text{VR}(p, S_\mathcal{P}) \cap D_\Gamma$ must be enclosed within any p -cycle of $\mathcal{J}_{s, S_\mathcal{P}} \cup \{\Gamma\}$, by its definition. Thus, no such p -cycle can be contained in $D_\mathcal{P}$. \square

Next, we give a key property of a Voronoi-like region $R(\alpha, \mathcal{P})$, called the *cut property*, see Fig. 11. Consider a connected component e of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ and let $\text{cut}(e)$ denote the portion of region $R(\alpha, \mathcal{P})$ that is cut out by e , as shown shaded in Fig. 11, and defined as follows. If e does not intersect α , let $\text{cut}(e)$ be the portion of the region at the opposite side of e as α , see Fig. 11a. If e is the only component of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ incident to α , let $\text{cut}(e)$ be the portion of $R(\alpha, \mathcal{P})$ incident to the side of e labeled s_β , see Fig. 11, b and d. If two different components of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ are incident to α , let $\text{cut}(e)$ be the portion of $R(\alpha, \mathcal{P})$ between these two components, see Fig. 11c. Note that if $\beta \in \mathcal{P}$ then only the cases (a) and (b) are possible. On the other hand, if $\mathcal{P} = \mathcal{E}$, and $\alpha, \beta \in \mathcal{E}$, then $J(s_\alpha, s_\beta)$ cannot intersect $\text{VR}(\alpha, \mathcal{E})$, thus, none of these cases is possible.

Lemma 3.8 Suppose bisector $J(s_\alpha, s_\beta)$ intersects $R(\alpha, \mathcal{P})$ (see Fig. 11). For any connected component e of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$, it holds $\text{cut}(e) \subseteq D(s_\beta, s_\alpha)$.

Proof Suppose first that a component e of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ does not intersect α , see Fig. 11a. Then the label s_α must appear on the same side of e as α , because otherwise, $\partial \text{cut}(e)$ would be an s_α -cycle, contradicting Lemma 3.7.

Suppose now that e intersects α . Then there is a component $\tilde{\beta}$ of $J(s, s_\beta) \cap R(\alpha, \mathcal{P})$, incident to the intersection point of e and α , that is contained in $\text{cut}(e)$. Since s -bisectors

can intersect at most twice, it follows that $\tilde{\beta}$ may have both its endpoints on α only if $\beta \notin \mathcal{P}$, because otherwise, $J(s, s_\beta)$ and $J(s, s_\alpha)$ would intersect more than twice. Thus, if $\beta \in \mathcal{P}$, e may only have one endpoint on α , and no other component of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ may be incident to α , see Fig. 11b. Otherwise, $J(s_\alpha, s_\beta)$ may intersect α twice, resulting in cases (c) or (d) of Fig. 11. No other cases exist.

Consider an arbitrary component e of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$. Suppose for the sake of contradiction that $\text{cut}(e) \not\subseteq D(s_\beta, s_\alpha)$. Then $J(s_\beta, s_\alpha)$ must intersect the interior of $\text{cut}(e)$ with a component e' of $J(s_\beta, s_\alpha) \cap R(\alpha, \mathcal{P})$, $e' \neq e$. Among any such component, let e' be the first one following e in the direction away from α . Since e' cannot intersect e nor can it intersect α , it follows that e' must create an s_α -cycle with $\partial\text{cut}(e)$, contradicting Lemma 3.7. Figure 17 illustrates such a forbidden s_γ -cycle created by a piece of $J(s_\beta, s_\gamma)$, shown in dashed lines, and $\partial R(\gamma, \mathcal{P})$. \square

Lemma 3.8 implies that any components of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ must appear sequentially along $\partial R(\alpha, \mathcal{P})$. That is, in a traversal of $\partial R(\alpha, \mathcal{P})$, starting at α , no component of $J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$ may appear between the endpoints of another. Further, if $J(s_\alpha, s_\beta)$ intersects $R(\alpha, \mathcal{P})$, then $J(s, s_\beta)$ must also intersect the domain $D_{\mathcal{P}}$. We use this fact to establish that $\mathcal{V}_l(\mathcal{P})$ is unique in the following theorem; the proof is deferred to Sect. 5.

Theorem 3.9 *Given a boundary curve \mathcal{P} of $S' \subseteq S$, $\mathcal{V}_l(\mathcal{P})$ is unique, assuming it exists.*

The complexity of $\mathcal{V}_l(\mathcal{P})$ is $O(|\mathcal{P}|)$ as it is a planar acyclic 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} on a set of core arcs $S' \subset S$ and its Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$. Let β^* be a core arc in $S \setminus S'$. We define an insertion operation \oplus , which adds β^* to \mathcal{P} , and derives the boundary curve $\mathcal{P}_\beta = \mathcal{P} \oplus \beta^*$ and its Voronoi-like diagram $\mathcal{V}_l(\mathcal{P}_\beta) = \mathcal{V}_l(\mathcal{P}) \oplus \beta^*$. Since β^* is a core arc, it must be entirely contained in the closure of the domain $D_{\mathcal{P}}$.

Given \mathcal{P} and β^* , let $\beta \supseteq \beta^*$ be the connected component of $J(s, s_\beta) \cap \overline{D_{\mathcal{P}}}$ that contains β^* (see Fig. 12). \mathcal{P}_β is the boundary curve derived from \mathcal{P} by substituting its portion between the endpoints of β , with β itself. We say that \mathcal{P}_β is derived from \mathcal{P} by *inserting* the core arc β^* , or equivalently, by inserting the original arc β . The insertion operation performs the following tasks algorithmically:

- Insert the core arc β^* in \mathcal{P} , deriving $\mathcal{P}_\beta = \mathcal{P} \oplus \beta^* = \mathcal{P} \oplus \beta$. The various cases are illustrated in Fig. 13, see Observation 4.1 below.
- Compute the *merge curve* $J(\beta)$, which defines the boundary of $R(\beta, \mathcal{P}_\beta)$.
- Update $\mathcal{V}_l(\mathcal{P})$, by inserting $J(\beta)$ and deleting any portion of the diagram enclosed by it, to derive $\mathcal{V}_l(\mathcal{P}_\beta) = \mathcal{V}_l(\mathcal{P}) \oplus \beta$.

These tasks are standard in relation to site insertion in any Voronoi diagram. We prove their correctness in a Voronoi-like structure, see Theorems 4.3 and 4.4.




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




Fig. 13 Insertion cases for an arc β

Observation 4.1 All possible cases of inserting arc $\beta^* \subseteq \beta$ in \mathcal{P} are enumerated as follows (see Fig. 13).

- (a) Arc 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 ; $R(\alpha, \mathcal{P})$ will also be split in two parts.
- (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 $\tilde{\beta} \supseteq \beta^*$; then $\beta = \tilde{\beta}$ and $\mathcal{P}_\beta = \mathcal{P}$.

In terms of auxiliary arcs, \mathcal{P}_β may contain fewer, the same number, or even one additional auxiliary arc as compared to \mathcal{P} .

Given $\mathcal{V}_l(\mathcal{P})$ and arc β , we define a merge curve $J(\beta)$, which delimits the boundary of $R(\beta, \mathcal{P}_\beta)$. We define $J(\beta)$ algorithmically (see Def. 4.2), starting at an endpoint of β , and tracing s_β -related bisectors within the faces of $\mathcal{V}_l(\mathcal{P})$, refer to Fig. 14. We prove that $J(\beta)$ is indeed an s_β -monotone path that connects the endpoints of β (Theorem 4.3). Let x, y denote the endpoints of β , where $x\beta y$ appear in counterclockwise order. We assume a counterclockwise traversal of \mathcal{P} . Refer to Fig. 14.

Definition 4.2 Given $\mathcal{V}_l(\mathcal{P})$ and arc $\beta \subseteq J(s, s_\beta)$, the merge curve $J(\beta)$ is a path (v_1, \dots, v_m) in the arrangement of s_β -related bisectors, $\mathcal{A}(\mathcal{J}_{s_\beta, s_\mathcal{P}} \cup \{\Gamma\})$, 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 a bisector edge, or an arc on Γ . We assume a clockwise ordering of $J(\beta)$. 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.

- (i) If $e_i \subseteq J(s_\beta, s_\alpha)$, let v_{i+1} be the other endpoint of the connected component of $J(s_\beta, s_\alpha) \cap R(\alpha, \mathcal{P})$ 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 Fig. 14, see $e_i = e', v_i = z, v_{i+1} = z'$.)
- (ii) If $e_i \subseteq \Gamma$, let g be the Γ -arc in \mathcal{P} incident to v_i , in clockwise order. Let $e_{i+1} \subseteq J(s_\beta, s_\gamma)$, where $\gamma \in \mathcal{P}$ and $R(\gamma, \mathcal{P})$ is the first region, incident to g clockwise




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




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




Fig. 16 J_x^i and J_y^j in Sect. 4.1

from v_i such that $J(s_\beta, s_\gamma)$ intersects $g \cap \overline{R(\gamma, \mathcal{P})}$; let v_{i+1} be this intersection point. (In Fig. 14, see $v_i = v$ and $v_{i+1} = w$.)

The following theorem shows that $J(\beta)$ forms an s_β -monotone path joining the endpoints of β . We defer its proof to the end of this section (Sect. 4.1).

Theorem 4.3 *The merge curve $J(\beta)$ is a unique s_β -monotone path in the arrangement of s_β -related bisectors $\mathcal{A}(\mathcal{I}_{s_\beta, s_\mathcal{P}} \cup \Gamma)$ connecting the endpoints of β . Further:*

- If arc β splits a single arc $\alpha \in \mathcal{P}$ (case (c) of Observation 4.1) then $J(\beta)$ must intersect $R(\alpha, \mathcal{P})$ in two different components, $e_1, e_{m-1} \subseteq J(s_\alpha, s_\beta)$. $J(\beta)$ can intersect any other region in $\mathcal{V}_l(\mathcal{P})$ at most once.
- $J(\beta)$ cannot intersect the region of any arc in $\mathcal{P} \setminus \mathcal{P}_\beta$, which gets deleted by the insertion of β , nor can it intersect arc β in its interior.

Let $T(\beta)$ denote the portion of $\mathcal{V}_l(\mathcal{P})$ enclosed by $J(\beta)$ and $\mathcal{P} \setminus \mathcal{P}_\beta$. Let $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ denote the graph obtained from $\mathcal{V}_l(\mathcal{P})$ by deleting $T(\beta)$ and substituting it with $J(\beta)$, i.e., $\mathcal{V}_l(\mathcal{P}) \oplus \beta = (\mathcal{V}_l(\mathcal{P}) \setminus T(\beta)) \cup J(\beta)$,




Fig. 17 Impossible configuration of $J(s_\beta, s_\gamma)$. Scanning $\partial R(\gamma, \mathcal{P})$ from v_i counterclockwise, Lemma 3.8 assures that v_{i+1} is the first encountered intersection of $J(s_\beta, s_\gamma)$ with $\partial R(\gamma, \mathcal{P})$

Theorem 4.4 $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ is the Voronoi-like diagram $\mathcal{V}_l(\mathcal{P}_\beta)$.

Proof By construction, $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ induces a subdivision of the domain $D_{\mathcal{P}_\beta}$. By Theorem 4.3, $J(\beta)$, and thus, $\partial R(\beta) \setminus \beta$, is an s_β -monotone path connecting the endpoints of β . For any arc $\alpha \in \mathcal{P}$ such that $J(\beta)$ passes through $R(\alpha, \mathcal{P})$, the boundary of the updated face in $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ remains an s_α -monotone path, by the definition of $J(\beta)$. Thus, for any face f of $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ incident to an arc $\alpha \neq \beta$, its boundary $\partial f \setminus \alpha$ is an s_α -monotone path, hence, it satisfies the first requirement of Definition 3.3.

Since $J(\beta)$ can enter any region in $\mathcal{V}_l(\mathcal{P})$ at most once (except from case (c) of Observation 4.1) it cannot create a face that may remain in the interior of $D_{\mathcal{P}}$. Further, $J(\beta)$ cannot pass through any region of an arc in $\mathcal{P} \setminus \mathcal{P}_\beta$, thus, such a region must be enclosed by $J(\beta)$ and will be deleted. Hence, any face of $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ must be incident to a boundary arc of \mathcal{P}_β , satisfying also the second requirement of Definition 3.3. Since, by Theorem 3.9, the Voronoi-like diagram of a boundary curve is unique, it follows that $\mathcal{V}_l(\mathcal{P}) \oplus \beta = \mathcal{V}_l(\mathcal{P}_\beta)$. \square

The tracing of the merge curve $J(\beta)$ within $\mathcal{V}_l(\mathcal{P})$ can be performed similarly to any ordinary Voronoi diagram (see, e.g., [2, Ch. 7.5.3]). This is correct in a Voronoi-like diagram as a result of the cut property of Lemma 3.8: when $J(\beta)$ enters a region $R(\gamma, \mathcal{P})$ at a point v_i , we can determine v_{i+1} by scanning $\partial R(\gamma, \mathcal{P})$ counterclockwise sequentially, until we encounter the first intersection with $J(s_\beta, s_\gamma)$. Lemma 3.8 assures that no intersection of $J(s_\beta, s_\gamma)$ with $\partial R(\gamma, \mathcal{P})$ between v_i and v_{i+1} is possible, such as the one shown in Fig. 17. Thus, we can state the following fact.

Lemma 4.5 Let $e_i = (v_i, v_{i+1})$ be an edge of $J(\beta)$ in $R(\gamma, \mathcal{P})$. Given v_i , we can determine v_{i+1} by sequentially scanning $\partial R(\gamma, \mathcal{P})$ counterclockwise from v_i (i.e., away from γ) until the first intersection of $J(s_\beta, s_\gamma)$ with $\partial R(\gamma, \mathcal{P})$ which determines v_{i+1} .

Special care is required in cases (c), (d), and (e) of Observation 4.1 to identify the first edge of $J(\beta)$, as β does not overlap any feature of $\mathcal{V}_l(\mathcal{P})$ in these cases. To handle them we need to define some additional parameters.

Let $\tilde{\mathcal{P}}$ denote the finer version of \mathcal{P} derived by intersecting its Γ -arcs with $\mathcal{V}_l(\mathcal{P})$, i.e., partitioning the Γ -arcs of \mathcal{P} into finer pieces by the incident faces of $\mathcal{V}_l(\mathcal{P})$. Since the complexity of $\mathcal{V}_l(\mathcal{P})$ is $O(|\mathcal{P}|)$, it follows that $|\tilde{\mathcal{P}}|$ is also $O(|\mathcal{P}|)$.

Definition 4.6 Let α and γ denote the original arcs preceding and following β on \mathcal{P}_β . We assume a counterclockwise traversal of \mathcal{P} and \mathcal{P}_β .

- (i) Let $d_1(\beta, \mathcal{P}_\beta)$ denote the number of auxiliary arcs that appear on \mathcal{P}_β from α to β .
- (ii) Let $d_2(\beta, \mathcal{P}_\beta)$ denote the number of auxiliary arcs that appear on \mathcal{P} between the endpoints of β that get deleted by the insertion of β .
- (iii) In case (c) of Observation 4.1, where β splits an arc ω in two arcs (ω_1, ω_2) , let $r(\beta, \mathcal{P}_\beta) = \min \{|\partial R(\omega_1, \mathcal{P}_\beta)|, |\partial R(\omega_2, \mathcal{P}_\beta)|\}$; in other cases, let $r(\beta, \mathcal{P}_\beta) = 0$.
- (iv) In case (d) of Observation 4.1, where β splits a Γ -arc, let $\tilde{d}(\beta, \mathcal{P}_\beta)$ denote the number of fine Γ -arcs on $\tilde{\mathcal{P}}_\beta$ from α to β (i.e., the number of regions in $\mathcal{V}_l(\mathcal{P}_\beta)$ incident to Γ from α to β); in all other cases, $\tilde{d}(\beta, \mathcal{P}_\beta) = 0$.

Lemma 4.7 *Given α, γ , and $\mathcal{V}_l(\mathcal{P})$, the merge curve $J(\beta)$ can be computed in time $O(|J(\beta)| + d_1(\beta, \mathcal{P}_\beta) + d_2(\beta, \mathcal{P}_\beta) + r(\beta, \mathcal{P}_\beta) + \tilde{d}(\beta, \mathcal{P}_\beta))$.*

Proof We assume a counterclockwise (ccw) ordering of \mathcal{P} . We first determine the endpoints of β in time $O(d_1(\beta, \mathcal{P}_\beta) + d_2(\beta, \mathcal{P}_\beta))$ by scanning sequentially the arcs in \mathcal{P} starting at α and moving ccw (towards γ) until the endpoints of β are determined. Note that β contains the core arc β^* , therefore, we can easily identify the correct component of $J(s, s_\beta) \cap D_{\mathcal{P}}$ during the scan, even if $J(s, s_\beta)$ intersects \mathcal{P} multiple times. This scan also determines which case of Observation 4.1 is relevant.

Let $T(\beta)$ denote the portion of $\mathcal{V}_l(\mathcal{P})$ that is enclosed by $J(\beta)$ and $\mathcal{P} \setminus \mathcal{P}_\beta$. $T(\beta)$ gets deleted by the insertion of β . It is an embedded forest, which by Theorem 4.3 is incident to the following faces of $\mathcal{V}_l(\mathcal{P})$: one face for each bisector edge of $J(\beta)$, and one face for each auxiliary arc $\alpha' \in \mathcal{P} \setminus \mathcal{P}_\beta$. The latter number is counted in $d_2(\beta, \mathcal{P}_\beta)$. We infer that $T(\beta)$ has complexity $O(|J(\beta)| + d_2(\beta, \mathcal{P}_\beta))$.

To compute $J(\beta)$, we trace $T(\beta)$ in time $O(|T(\beta)|)$, after having identified one of its leaves, as normally done in an ordinary Voronoi diagram. This statement is correct due to Theorem 4.3 and Lemma 4.5. However, we first need to identify one leaf of $T(\beta)$, and certain cases of Observation 4.1 may require additional scans, which can increase the time complexity over $|T(\beta)|$. We give the case analysis in the remainder of this proof.

Suppose first that $T(\beta)$ has a leaf on \mathcal{P} . Then, in all cases of Observation 4.1, except cases (d) and (e), a leaf of $T(\beta)$ is identified by the initial scan. In case (e), β has at least one endpoint on a boundary arc ρ of \mathcal{P} , see Fig. 14; we identify a leaf by scanning $\tilde{\mathcal{P}}$ starting at ρ and moving towards the other endpoint of β . This scan takes only one step as the leaf will be incident to the first Γ -arc neighboring ρ on $\tilde{\mathcal{P}}$. In case (d), both endpoints of β are on Γ . We scan $\tilde{\mathcal{P}}$ from α to β until we locate the first endpoint x of β . A leaf of $T(\beta)$ must be incident to the fine Γ -arc that contains x . Since all the encountered Γ -arcs remain in $\tilde{\mathcal{P}}_\beta$, the term $O(\tilde{d}(\beta, \mathcal{P}_\beta))$ is added to the overall time complexity.

Suppose now that $T(\beta)$ has no leaf on \mathcal{P} . Then β is enclosed within a single Voronoi-like region $R(\omega, \mathcal{P})$. There are three cases to consider: Observation 4.1, (c), (d), and (e).

In case Observation 4.1 (c), the insertion of β splits arc ω in two parts, ω_1 and ω_2 . We scan $\partial R(\omega, \mathcal{P})$ sequentially until an intersection with $J(s_\omega, s_\beta)$ is found. This intersection point is a leaf of $T(\beta)$ within the domain of \mathcal{P} . We start scanning from both endpoints of ω , tracing the shorter among $\partial R(\omega_1, \mathcal{P}_\beta)$ and $\partial R(\omega_2, \mathcal{P}_\beta)$. This adds the term $r(\beta, \mathcal{P}_\beta)$ to the overall time complexity.




Fig. 18 Case (e) of Observation 4.1, where $T(\beta)$ has no leaf on \mathcal{P} . Endpoint x lies on a fine Γ -arc g_1 bounding $R(\omega, \mathcal{P})$, and $y \in \omega$




Fig. 19 Case (d) of Observation 4.1, where $T(\beta)$ has no leaf on \mathcal{P} . Both x, y lie on a fine Γ -arc g_1 bounding $R(\omega, \mathcal{P})$

In cases (d) and (e) of Observation 4.1, $J(\beta) \subseteq R(\omega, \mathcal{P}) \cup \Gamma$, since otherwise $J(\beta)$ would intersect the region $R(\omega, \mathcal{P})$ twice, contradicting Theorem 4.3. Thus, $J(\beta)$ consists of a single bisector $J(s_\omega, s_\beta)$ and one or two Γ -arcs, see Figs. 18 and 19, respectively. Therefore, we only need to identify ω . In case (e), ω is identified during the initial scan. In case (d), β has both its endpoints on Γ , and we scan $\tilde{\mathcal{P}}$ from α to β until we encounter the fine Γ -arc that contains the first endpoint of β ; the latter Γ -arc bounds the region $R(\omega, \mathcal{P})$. This scan adds the term $O(\tilde{d}(\beta, \mathcal{P}_\beta))$ to the time complexity. \square

4.1 Proving Theorem 4.3

In this section we prove Theorem 4.3. The proof is technical but it is self-contained and it is not necessary for following the rest of the paper. We first establish the following lemma.

Lemma 4.8 *The merge curve $J(\beta)$ cannot intersect arc β , other than its endpoints.*

Proof Suppose that an edge e_i of $J(\beta)$, such that $e_i \subseteq J(s_\alpha, s_\beta)$ and $e_i \subseteq R(\alpha, \mathcal{P})$, intersects arc β . Then $J(s, s_\alpha)$ must also pass through the same intersection point within $R(\alpha, \mathcal{P})$. But an s -related bisector $J(s, s_\alpha)$ can never intersect $R(\alpha, \mathcal{P})$, by Lemma 3.6. \square

The following observation is used throughout the proofs in this section.

Lemma 4.9 *For any site $p \in S \setminus \{s\}$, $D(s, p) \cap D_{\mathcal{P}}$ is connected. Thus, any components of the same s -related bisector $J(s, \cdot) \cap D_{\mathcal{P}}$ must appear along \mathcal{P} sequentially, one after another.*

Proof If we assume the contrary, we obtain a forbidden s -inverse cycle defined by $J(s, \cdot)$ and \mathcal{P} , which contradicts Lemma 2.3. \square




Fig. 20 Illustrations for Lemma 4.10

We now establish that $J(\beta)$ cannot pass through any region of an auxiliary arc in $\mathcal{P} \setminus \mathcal{P}_\beta$ that gets deleted by the insertion of β .

Lemma 4.10 *Let $\alpha \in \mathcal{P}$ but $\alpha \not\subseteq \mathcal{P}_\beta$. Then $R(\alpha, \mathcal{P}) \subset D(s_\beta, s_\alpha)$, see Fig. 20.*

Proof By Lemma 3.6, it holds that $R(\alpha, \mathcal{P}) \subseteq D(s, s_\alpha)$. Let $R_s = R(\alpha, \mathcal{P}) \cap D(s, s_\beta)$ and $R_\beta = R(\alpha, \mathcal{P}) \cap D(s_\beta, s)$. By transitivity of dominance regions we have $R_\beta \subseteq D(s_\beta, s_\alpha)$. By Lemma 4.9, R_s is not incident to α . Thus, if $J(s_\beta, s_\alpha)$ intersected R_s then it would create an s_α -cycle with the boundary of $R(\alpha, \mathcal{P})$, contradicting Lemma 3.7, see the dashed gray line in Fig. 20. This also implies that $R_s \subseteq D(s_\beta, s_\alpha)$. Thus, $R(\alpha, \mathcal{P}) = R_s \cup R_\beta \subseteq D(s_\beta, s_\alpha)$. \square

In the following we prove that $J(\beta)$ is an s_β -monotone path connecting the endpoints of β . To this aim we perform a bi-directional induction on the vertices of $J(\beta)$.

Let $J_x^i = (v_1, v_2, \dots, v_i)$, $1 \leq i < m$, 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 Fig. 16. Note that vertex v_i uniquely determines e_i , however, its other endpoint is not yet specified. Similarly, let $J_y^j = (v_m, v_{m-1}, \dots, v_{m-j+1})$, $1 \leq j < m$, 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} . For any bisector edge $e_\ell \in J(\beta)$, let α_ℓ denote the boundary arc that induces e_ℓ , i.e., $e_\ell \subseteq J(s_{\alpha_\ell}, s_\beta) \cap R(\alpha_\ell, \mathcal{P})$.

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

Inductive step: Assuming that $i + j < m$, we prove that at least one of J_x^i or J_y^j can grow to J_x^{i+1} or J_y^{j+1} respectively at a *valid* vertex (Lemmas 4.11, 4.12), entering a new region of $\mathcal{V}_l(\mathcal{P})$ that has not been visited by J_x^i or J_y^j (Lemma 4.14). A vertex is called *valid* if it belongs to $\mathcal{A}(\mathcal{J}_{s_\beta, s_\mathcal{P}} \cup \{\Gamma\})$ or it is an endpoint of β . When $i + j = m$, a finish condition is given in Lemma 4.13. The base case for $i = j = 1$ is trivially true. In the remaining section we prove correctness of the inductive step.

Suppose that $e_i \subseteq J(s_{\alpha_i}, s_\beta)$ and $v_i \in \partial R(\alpha_i, \mathcal{P})$. To show that v_{i+1} is a valid vertex it is enough to show that (1) v_{i+1} cannot 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 lemmata we assume that the inductive hypothesis holds.




Fig. 21 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 4.11

Lemma 4.11 *Suppose $e_i \subseteq J(s_{\alpha_i}, s_\beta)$ but $v_{i+1} \in \alpha_i$, that is, e_i hits arc $\alpha_i \in \mathcal{P}$, and thus, v_{i+1} is not a valid vertex. Then vertex v_{m-j} must be a valid vertex in $\mathcal{A}(\mathcal{J}_{s_\beta, s_\mathcal{P}})$, and v_{m-j} cannot be on \mathcal{P} .*

Proof Suppose vertex v_{i+1} of e_i lies on arc α_i as shown in Fig. 21a. 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 $J(s, s_\beta)$. Thus, $v_1, v_m, v_{i+1} \in J(s, s_\beta)$. By the inductive hypothesis, no other vertex of J_x^i nor J_y^j can be on $J(s, s_\beta)$. Vertices v_1, v_{i+1}, v_m appear on \mathcal{P} in clockwise order, because J_x^{i+1} cannot intersect β . 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} must lie on J_2 , as otherwise, 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 Fig. 21a, contradicting Lemma 2.3. This cycle must be s_β -inverse because $J_x^{i+1} \subseteq \overline{D\mathcal{P}}$, and all components of $J(s, \cdot) \cap D\mathcal{P}$ must appear sequentially along \mathcal{P} by Lemma 4.9.

Thus, v_{i+1} lies on J_2 . Further, by Lemma 4.9, 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 Fig. 21b, which illustrates $J(s, s_\beta)$ 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} cannot 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 by the inductive hypothesis, α_{m-j} is distinct from all $\alpha_\ell, \ell \leq i$. 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 Fig. 21b, similarly to the first paragraph of this proof. 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 Fig. 21c. 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 $\mathcal{A}(\mathcal{J}_{s_\beta, s_\mathcal{P}})$. \square

The proof for the following lemma is similar.

Lemma 4.12 *Suppose vertex v_i is on a Γ -arc $g \in \mathcal{P}$ but v_{i+1} cannot be determined because no bisector $J(s_\beta, s_\gamma)$ intersects $\overline{R(\gamma, \mathcal{P})} \cap g$, 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} cannot be on \mathcal{P} .*




Fig. 22 The assumption that $v_i \in \Gamma$ and v_{i+1} of the merge curve J_x^i cannot be determined as in Lemma 4.12

Proof We truncate the Γ -arc g to its portion clockwise from v_i ; let w be the endpoint of g clockwise from v_i , see Fig. 22a. If no $J(s_\beta, s_\gamma) \cap R(\gamma, \mathcal{P})$ intersects g , as we assume in this lemma, then $R(\gamma, \mathcal{P}) \cap g \subseteq D(s_\beta, s_\gamma)$, for any region $R(\gamma, \mathcal{P})$ incident to g . Thus, $w \in D(s_\beta, s)$. However, $v_i \in D(s, s_\beta)$, since, by Lemma 3.6, $R(\alpha_{i-1}) \subseteq D(s, s_{\alpha_{i-1}})$ and v_i is incident to $J(s_\beta, s_{\alpha_{i-1}}) \cap R(\alpha_{i-1})$. Thus, $J(s, s_\beta)$ must intersect g at some point z clockwise from v_i . Arc β partitions $J(s, s_\beta)$ in two parts: J_1 incident to v_1 and J_2 incident to v_m . Lemma 4.9 implies that all components of $J_2 \cap D\mathcal{P}$ appear on \mathcal{P} clockwise after v_i and before v_m , as shown by the black dashed curve in Fig. 22a; also z lies on J_2 .

Now we can show that vertex v_{m-j} of J_y^j cannot be on \mathcal{P} analogously to the proof of Lemma 4.11. The only difference is that we must additionally show that v_{m-j} cannot lie on \mathcal{P} clockwise after v_i and before w . But this holds already by the assumption in the lemma statement. Refer to Fig. 22, b and c. We conclude that v_{m-j} cannot lie on \mathcal{P} and it is a valid vertex of $\mathcal{A}(J_{s_\beta, s_\mathcal{P}})$. \square

Lemma 4.13 in the sequel provides a finish condition for the induction, when J_x^i and J_y^j are incident to a common region or to a common Γ -arc. When it is met, the merge curve $J(\beta)$ is a concatenation of J_x^i and J_y^j .

Lemma 4.13 *Suppose $i + j > 2$ and either (1) or (2) holds: (1) v_i and v_{m-j+1} are incident to the same region $R(\alpha_i, \mathcal{P})$ and $e_i, e_{m-j} \subseteq J(s_\beta, s_{\alpha_i})$, i.e., $\alpha_i = \alpha_{m-j}$; or (2) v_i and v_{m-j+1} are on the same 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 $m = i + j$.*

Proof Let $\alpha = \alpha_i$. Suppose (1) holds, then $e_i, e_{m-j} \subseteq J(s_\beta, s_\alpha)$, see Fig. 23a. The boundary $\partial R(\alpha_i, \mathcal{P})$ is partitioned in four parts, using a counterclockwise traversal starting at α_i : 1. ∂R_1 , from the endpoint of arc α_i to v_i ; 2. ∂R_2 , from v_i to v_{m-j+1} ; 3. ∂R_3 , from v_{m-j+1} to the next endpoint of α_i ; and 4. arc α_i . We show that e_i and e_{m-j} cannot hit any of these parts, thus, $e_i = e_{m-j}$.

- (i) Edge e_i cannot hit ∂R_1 and edge e_{m-j} cannot hit ∂R_3 , by the cut property of Lemma 3.8.
- (ii) We prove that edge e_i cannot hit ∂R_2 (analogously for edge e_{m-j}). Let ρ be any edge on ∂R_2 . (If $v_i \in \rho$ or $v_{m-j+1} \in \rho$, assume that ρ is truncated with endpoint v_i or v_{m-j+1} respectively).




Fig. 23 Illustrations for Lemma 4.13. **a** corresponds to condition (1) and **b** to condition (2). The label $R(\gamma)$ abbreviates $R(\gamma, \mathcal{P})$ and the label $R(\alpha_i)$ abbreviates $R(\alpha_i, \mathcal{P})$

- Suppose that ρ is a bisector edge, $\rho \subseteq J(s_\alpha, s_\gamma)$, see Fig. 23a. Then at least one of J_y^j , J_x^i , or β must pass through $R(\gamma, \mathcal{P})$. Suppose that J_y^j does, as shown in Fig. 23a. Then, by the cut property (Lemma 3.8), $\rho \subseteq D(s_\beta, s_\gamma)$. By transitivity (Lemma 2.4) it also holds that $\rho \subseteq D(s_\beta, s_\alpha)$. Thus, e_i cannot hit ρ . Symmetrically for J_x^i . If only β passes through $R(\gamma, \mathcal{P})$, then we can use Lemma 4.10 to derive that $\rho \subseteq D(s_\beta, s_\gamma)$; the rest follows.
- Suppose that $\rho \subseteq \Gamma$. Then either ρ itself is part of an edge of J_y^j or of J_x^i , or β passes through $R(\alpha, \mathcal{P})$ and ρ is at opposite side of it than α . In the former case, $\rho \subseteq D(s_\beta, s_\alpha)$ by the definition of a Γ -edge in the merge curve. In the latter case, the same is derived by Lemma 3.6 and transitivity (Lemma 2.4). Thus, e_i cannot hit ρ .

- (iii) Edge e_i (resp. e_{m-j}) cannot hit ∂R_3 , because if it did, e_i and e_{m-j} would not appear sequentially on $R(\alpha_i, \mathcal{P})$ contradicting Lemma 3.8.
- (iv) It remains to show that e_i and e_{m-j} cannot both hit α_i ; however, this is already shown in Lemma 4.11.

Suppose now that (2) holds, see Fig. 23b. Let $R(\gamma, \mathcal{P})$ be a region in $\mathcal{V}_l(\mathcal{P})$ incident to the Γ -arc g and let $\rho = R(\gamma, \mathcal{P}) \cap g$ be the Γ -arc bounding $R(\gamma, \mathcal{P})$, which lies between v_i and v_{m-j+1} . At least one of J_y^j or J_x^i or β must pass through $R(\gamma, \mathcal{P})$. By the exact same arguments as before, $\rho \subseteq D(s_\beta, s_\gamma)$. We infer that there is no bisector $J(s_\beta, s_\gamma)$ in $R(\gamma, \mathcal{P})$, for any region $R(\gamma, \mathcal{P})$ incident to g between v_i and v_{m-j+1} . Thus, $e_{i+1} = e_{m-j+1} \subseteq g$.

We conclude that in both (1) and (2), $v_{i+1} = v_{m-j+1}$, $v_{m-j} = v_i$, and $m = i + j$. $J(\beta)$ is the concatenation of J_x^i and J_y^j with $e_{i+1} = e_{m-j+1}$. □

Lemma 4.14 *Suppose vertex v_{i+1} is valid and $e_{i+1} \subseteq J(s_\beta, s_{\alpha_{i+1}})$. Then $R(\alpha_{i+1}, \mathcal{P})$ 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 $m - j < \ell$.*

Proof Let $e_k, k \leq i$, be a bisector edge of J_x^i . Denote by ∂R_k^1 the portion of $\partial R(\alpha_k, \mathcal{P})$ from α_k to v_k in a counterclockwise traversal, see the bold red part ∂R_i^1 in Fig. 24.




Fig. 24 Illustration for Lemma 4.14

Analogously, for a bisector edge e_{m-j} of J_y^j , where ∂R_{m-j}^1 is defined in a clockwise traversal of $\partial R(\alpha_{m-j}, \mathcal{P})$. Recall that $\text{cut}(e_k)$ denotes the portion of $R(\alpha_k, \mathcal{P})$ cut out by edge e_k , at opposite side from α_k .

The cut property of Lemma 3.8 implies that v_{i+1} cannot be on $\partial \text{cut}(e_\ell)$ for any ℓ , $\ell < i$ and $m - j < \ell$, and that v_{i+1} cannot be on ∂R_i^1 . This implies that v_{i+1} cannot be on ∂R_ℓ^1 for any $\ell < i$, because we have a plane graph in $D_{\mathcal{P}}$ and by its layout ∂R_ℓ^1 is not reachable from e_i without first hitting $\partial \text{cut}(e_\ell)$ or ∂R_i^1 . See Fig. 24. Thus, v_{i+1} cannot be on $\partial R(\alpha_\ell)$, $\ell < i$. By Lemma 4.13, v_{i+1} cannot be on ∂R_{m-j}^1 . This implies, again by the layout, that v_{i+1} cannot be on ∂R_ℓ^1 for all $\ell > m - j$. Thus, v_{i+1} cannot be on $\partial R(\alpha_\ell, \mathcal{P})$, for any $\ell > m - j$. This implies that $\alpha_{i+1} \neq \alpha_\ell$, for any ℓ , $\ell \leq i$ or $\ell > m - j$. \square

By Lemma 4.14, 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 by a lower index edge. Hence, conditions (1) or (2) of Lemma 4.13 must be fulfilled at some point of the induction, completing the proof of Theorem 4.3.

Completing the bi-directional induction establishes also the remaining properties for $J(\beta)$. First, $J(\beta)$ can never enter the same region twice (by Lemma 4.14), except the region of α_1 , if $\alpha_1 = \alpha_m$. The latter is Observation 4.1(c), where arc β splits a single arc $\alpha \in \mathcal{P}$. In this case $J(\beta)$ enters $R(\alpha, \mathcal{P})$ exactly twice and both $e_1, e_{m-1} \subseteq J(s_\alpha, s_\beta)$. This is because $J(\beta)$ must intersect $\partial R(\alpha, \mathcal{P})$, i.e., $J(\beta) \not\subseteq R(\alpha, \mathcal{P})$, as otherwise $J(\beta) = J(s_\alpha, s_\beta)$ (see Fig. 15) contradicting the labeling of the cut property in Lemma 3.8.

Completing the induction for Theorem 4.3 establishes also that $J(\beta)$ is unique and that the conditions of Lemmas 4.11 and 4.12 can never be met. Thus, no vertex of $J(\beta)$, except its endpoints, can be on a boundary arc of \mathcal{P} .

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

In this section we prove Theorem 3.9 and establish that the Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$ is unique, for any boundary curve \mathcal{P} on $S' \subseteq S$. We first use Theorem 4.3 to show an essential property of Voronoi-like regions, which completes and extends the cut property of Lemma 3.8.

Lemma 5.1 *Let \mathcal{P} be a boundary curve on S' , $\mathcal{P} \neq \mathcal{E}$, and let $\alpha, \beta \in \mathcal{P}$ be two arcs such that $s_\alpha \neq s_\beta$. Suppose that $J(s_\alpha, s_\beta)$ intersects $R(\alpha, \mathcal{P})$ with a component e ,*




Fig. 25 A component e of $J(s_\alpha, \cdot)$ in $R(\alpha, \mathcal{P})$ as in Lemma 5.1




Fig. 26 A component e of $J(s_\alpha, \cdot)$ in $R(\alpha, \mathcal{P})$ with its endpoint v on a Γ -arc g as in Lemma 5.1

$e \subseteq J(s_\alpha, s_\beta) \cap R(\alpha, \mathcal{P})$. Then, $J(s, s_\beta)$ must intersect the domain $D_{\mathcal{P}}$. Further, there exists a component β' of $J(s, s_\beta) \cap D_{\mathcal{P}}$ such that the merge curve $J(\beta')$ in $\mathcal{V}_l(\mathcal{P})$ contains e , i.e., $e \subseteq \partial R(\beta', \mathcal{P} \oplus \beta')$.

We say that the arc β' is missing from \mathcal{P} .

Proof Suppose that a component e of $J(s_\alpha, s_\beta)$ intersects $R(\alpha, \mathcal{P})$, however, $J(s, s_\beta)$ does not intersect $D_{\mathcal{P}}$, i.e., $D_{\mathcal{P}} \subseteq D(s, s_\beta)$. Then, for any arc $\chi \in \mathcal{P}$, $\chi \subseteq \mathcal{J}(s, s_\chi)$ and $\chi \subseteq D(s_\chi, s_\beta)$, by the transitivity of dominance regions (Lemma 2.4). Let $\text{cut}(e)$ denote the portion of $R(\alpha, \mathcal{P})$ cut out by e , at opposite side from α , as defined in Lemma 3.8; then $\text{cut}(e) \subseteq D(s_\beta, s_\alpha)$, by Lemma 3.8.

Consider an endpoint v of e . There are two cases:

- (i) If v is on an edge ρ incident to regions $R(\alpha, \mathcal{P})$ and $R(\gamma, \mathcal{P})$, then $J(s_\beta, s_\gamma)$ intersects $R(\gamma, \mathcal{P})$ by an edge e_ρ , incident to v , leaving ρ and γ at opposite sides, since $D_{\mathcal{P}} \subseteq D(s, s_\beta)$, implying that $\gamma \subseteq D(s_\gamma, s_\beta)$, see Fig. 25.
- (ii) If v is on a Γ -arc g , let $R(\gamma, \mathcal{P})$ be the first region after v (on the side of e labeled s_β) such that $J(s_\beta, s_\gamma)$ intersects $g \cap \overline{R(\gamma, \mathcal{P})}$ at a point u (see Fig. 26). Such a region must exist because for all boundary arcs $\chi \in \mathcal{P}$, including the ones incident to g , $\chi \subseteq D(s_\chi, s_\beta)$. Let e_g be the component of $J(s_\beta, s_\gamma) \cap R(\gamma, \mathcal{P})$ incident to u .

Therefore, given e and v , we derive an edge e' , either $e' = e_\rho$ or $e' = e_g$, with the same properties as e , in a different region of $\mathcal{V}_l(\mathcal{P})$. This process repeats and there is no way to break it because for any arc $\chi \in \mathcal{P}$, $\chi \subseteq D(s_\chi, s_\beta)$. Thus, we create a closed curve on $\mathcal{V}_l(\mathcal{P})$ consisting of consecutive pieces of $J(s_\beta, \cdot)$, possibly interleaved with Γ -arcs, which has the label s_β in its interior. No two edges of this curve can intersect in their interior, within a region $R(\chi, \mathcal{P})$, because these edges would be pieces of the




Fig. 27 Arc $\beta' \subseteq J(s, s_\beta)$ in $D_{\mathcal{P}}$. The merge curve $J(\beta')$ contains e

same bisector $J(s_\beta, s_\chi)$, which in turn would not be a simple curve. For exactly the same reason, the curve may not enter $R(\beta, \mathcal{P})$. Furthermore, no vertex of the curve can repeat under our general position assumption, as no three s_β -related bisectors can intersect at the same point. Thus, the closed curve must be an s_β -cycle C that is contained in $D_{\mathcal{P}}$, see Fig. 25, which contradicts Lemma 3.7. Thus, our assumption that $J(s, s_\beta) \cap D_{\mathcal{P}} = \emptyset$ was false, and hence, $J(s, s_\beta)$ must intersect \mathcal{P} .

Let J_e denote the sequence of encountered edges e_ρ , starting with the initial edge e and ending on the first intersection of an arc χ_0 in \mathcal{P} with $J(s, s_\beta)$. Let β' be the component of $J(s, s_\beta) \cap D_{\mathcal{P}}$ incident to χ_0 , see Fig. 27. Clearly $\beta' \neq \beta$, as otherwise J_e would have entered $R(\beta, \mathcal{P})$. Consider the merge curve $J(\beta')$ for the arc β' on $\mathcal{V}_l(\mathcal{P})$ (see Definition 4.2). By its definition, the path J_e must be a portion of $J(\beta')$. Since by Theorem 4.3 the merge curve $J(\beta')$ on $\mathcal{V}_l(\mathcal{P})$ is unique, it follows that $J(\beta')$ contains J_e , and thus, it also contains edge e . \square

Note that no arc can be missing from the envelope \mathcal{E} of S' . We can now prove Theorem 3.9 from Sect. 3.

Theorem 3.9 *Given a boundary curve \mathcal{P} for $S' \subseteq \mathbb{S}$, $\mathcal{V}_l(\mathcal{P})$ is unique.*

Proof Let \mathcal{P} be a boundary curve for $S' \subseteq \mathbb{S}$ such that \mathcal{P} admits a Voronoi-like diagram $\mathcal{V}_l(\mathcal{P})$. Suppose there exist two different Voronoi-like diagrams of \mathcal{P} , $\mathcal{V}_l^{(1)} \neq \mathcal{V}_l^{(2)}$. Then there must be an edge $e^{(1)}$ of $\mathcal{V}_l^{(1)}$ bounding regions $R^{(1)}(\alpha, \mathcal{P})$ and $R^{(1)}(\beta, \mathcal{P})$ of $\mathcal{V}_l^{(1)}$, where $\alpha, \beta \in \mathcal{P}$, such that $e^{(1)}$ intersects region $R^{(2)}(\alpha, \mathcal{P})$ of $\mathcal{V}_l^{(2)}$, since α is common to both $R^{(1)}(\alpha, \mathcal{P})$ and $R^{(2)}(\alpha, \mathcal{P})$.

Let edge $e \subseteq J(s_\beta, s_\alpha)$ be the component of $R^{(2)}(\alpha, \mathcal{P}) \cap J(s_\beta, s_\alpha)$ overlapping with $e^{(1)}$, see Fig. 28. From Lemma 5.1, it follows that there is a non-empty component β_0 of $J(s, s_\beta) \cap D_{\mathcal{P}}$ such that $J(\beta_0)$ in $\mathcal{V}_l^{(2)}$ contains edge e . Since $J(\beta_0)$ and $\partial R^{(1)}(\beta, \mathcal{P})$ have an overlapping portion $e \cap e^{(1)}$ and they bound the regions of two different arcs $\beta_0 \neq \beta$ of site s_β , they form an s_β -cycle C as shown in Fig. 28. But C is contained in $D_{\mathcal{P}}$, deriving a contradiction to Lemma 3.7. \square




Fig. 28 Illustrations for the proof of Theorem 3.9

6 A Randomized Incremental Algorithm

Consider a random permutation $o = (\alpha_1, \dots, \alpha_h)$ of the set \mathcal{S} of core arcs, where $h = |\mathcal{S}|$. For $1 \leq i \leq h$, define the set $\mathcal{S}_i = \{\alpha_1, \dots, \alpha_i\} \subseteq \mathcal{S}$ to be the subset of the first i arcs in o , and permutation $o_i = (\alpha_1, \dots, \alpha_i)$. Let \mathcal{P}_i denote the boundary curve derived by the arc insertion operation \oplus by considering arcs in the order o_i . Let D_i denote the corresponding domain enclosed by \mathcal{P}_i .

Our randomized algorithm is inspired by the randomized, two-phase, approach of Chew [7] for the Voronoi diagram of points in convex position. Here the sites are core arcs in \mathcal{S} , forming boundary curves, and the algorithm constructs Voronoi-like diagrams within a series of shrinking domains $D_i \supseteq D_{i+1}$. The domain of \mathcal{P}_1 is $D_1 = D(s, s_{\alpha_1}) \cap D_\Gamma$; and D_h coincides with the Voronoi region $\text{VR}(s, S) \cap D_\Gamma$. The boundary curves are obtained by the insertion operation \oplus , one at each step, starting with $\mathcal{P}_1 = J(s, s_{\alpha_1}) \cap D_\Gamma$, and ending with $\mathcal{P}_h = \partial \text{VR}(s, S) \cap D_\Gamma$. The algorithm works in two phases.

In phase 1, the core arcs in \mathcal{S} get deleted one by one, in the reverse order of o , while recording the neighbors of an arc at the time of its deletion. Let $\mathcal{P}_1 = J(s, s_{\alpha_1}) \cap D_\Gamma$, $R(\alpha_1, \mathcal{P}_1) = D(s, s_{\alpha_1}) \cap D_\Gamma$, and $\mathcal{V}_l(\mathcal{P}_1) = \emptyset$.

In phase 2, we start with $\mathcal{V}_l(\mathcal{P}_1)$ and incrementally compute $\mathcal{V}_l(\mathcal{P}_i)$, $i = 2, \dots, h$, by inserting arc α_i to \mathcal{P}_{i-1} , and obtaining $\mathcal{P}_i = \mathcal{P}_{i-1} \oplus \alpha_i$, and $\mathcal{V}_l(\mathcal{P}_i) = \mathcal{V}_l(\mathcal{P}_{i-1}) \oplus \alpha_i$. When considering an arc α_i , we use the information of its recorded neighbors from phase 1 to determine its insertion point. At the end, we obtain $\mathcal{V}_l(\mathcal{P}_h)$, where \mathcal{P}_h is a boundary curve for \mathcal{S} . Since \mathcal{S} has only one boundary curve, it follows that \mathcal{P}_h coincides with $\partial \text{VR}(s, S) \cap D_\Gamma$.

We have already established the correctness of the insertion operation \oplus , thus, the algorithm correctly computes $\mathcal{V}_l(\mathcal{P}_h)$. We have also established that $\mathcal{V}_l(\mathcal{P}_h)$ coincides with the true Voronoi diagram $\mathcal{V}(\mathcal{S})$, by Corollary 3.5. Thus, the algorithm correctly computes $\mathcal{V}_l(\mathcal{P}_h) = \mathcal{V}(\mathcal{S}) = \mathcal{V}(\mathcal{S} \setminus \{s\}) \cap \text{VR}(s, S) \cap D_\Gamma$.

Next we analyze the time complexity of this algorithm and prove that the time complexity of step- i is expected $O(1)$. Thus, the overall time complexity is expected $O(h)$.

Lemma 6.1 \mathcal{P}_i contains at most $i - 1$ auxiliary arcs; thus, $|\mathcal{V}_l(\mathcal{P}_i)| = O(i)$.

Proof By definition, $|\mathcal{P}_1| = 1$. At each step of phase 2, exactly one original arc is inserted, and at most one additional auxiliary arc is created by a split in case (c) of




Fig. 29 There are $h!/(h - i)!$ nodes nodes at level- i of the decision tree \mathcal{T} , each corresponding to a unique permutation of i core arcs; the label of a node indicates the last element in the permutation. Level i is partitioned into disjoint groups of i nodes (permutations) each; $(i - 1)!$ such groups constitute a block Π_i . The illustration is schematic, the grouped nodes are not consecutive.

Observation 4.1, except from $i = 1$ and $i = h$. Thus, the total number of auxiliary arcs is at most $i - 1$ and the number of original arcs is at most i . Since an original arc may be merged with its neighbor in case (f) of Observation 4.1, the number of original arcs in \mathcal{P}_i may indeed be less than i . Since the complexity of $\mathcal{V}_i(\mathcal{P}_i)$ is $O(|\mathcal{P}_i|)$, the claim follows. \square

6.1 Time Analysis of the Randomized Incremental Algorithm, a Variant of Backwards Analysis

The time complexity of the algorithm, for each step i , has been expressed in Lemma 4.7 as a function of the resulting diagram $\mathcal{V}_i(\mathcal{P}_i)$. This calls for *backwards analysis* to estimate its expectation, see [20]. However, although $\mathcal{V}_i(\mathcal{P}_i)$ is unique, the boundary curve \mathcal{P}_i , and consequently its diagram, depend on the permutation order. As a result, backwards analysis is not directly applicable, contrary to our preliminary paper. In this section we revisit the analysis of [9], and introduce a variation of backwards analysis that is applicable to order-dependent structures.

Consider the *decision tree* \mathcal{T} of all possible random choices that can be made by our incremental algorithm on the input set of core arcs \mathcal{S} , $h = |\mathcal{S}|$, see Fig. 29. \mathcal{T} has $h!$ leaves each corresponding to one permutation of the arcs in \mathcal{S} . At level- i , there are $h!/(h - i)!$ nodes, and each node corresponds to a unique permutation of i core arcs. A set of i core arcs \mathcal{S}_i is associated with $i!$ different nodes at level- i , which are called the *block of \mathcal{S}_i* . We have $\binom{h}{i}$ distinct such blocks at level- i . Although all nodes within one block are associated with the same set of core arcs, their corresponding boundary curves may vary considerably depending on their permutation order.

We use the following strategy. We partition each block at level- i into $(i - 1)!$ disjoint groups of i nodes each. For each group we show that step i requires total time $O(i)$, considering all the i permutations within the group. Thus, on average, the algorithm spends $O(1)$ time on each node of \mathcal{T} . Since all permutations are equally likely, we obtain the expected linear $O(h)$ time complexity of our algorithm.




Fig. 30 Schematic differences between the boundary curves $\mathcal{B}_1, \dots, \mathcal{B}_i$. The domain D_i is shown shaded

Let $o_i = (\alpha_1, \alpha_2, \dots, \alpha_i)$ be an arbitrary permutation of \mathcal{S}_i . From o_i we define a group $G = G(o_i)$ of i permutations as follows: for each $1 \leq j < i$, remove α_j from its position in o_i and append it to the end of o_i .

$$o_i = (\alpha_1, \alpha_2, \dots, \alpha_{j-1}, \boxed{\alpha_j}, \alpha_{j+1}, \dots, \alpha_{i-1}, \alpha_i), \tag{1}$$

$$o_j = (\alpha_1, \alpha_2, \dots, \alpha_{j-1}, \alpha_{j+1}, \dots, \alpha_{i-1}, \alpha_i, \boxed{\alpha_j}) \tag{2}$$

Let \mathcal{B}_j and $\mathcal{V}_l(\mathcal{B}_j)$, $1 \leq j \leq i$, denote the boundary curves, and their Voronoi-like diagrams, derived incrementally, by arc insertion, following the order o_j , see Fig. 30. The boundary curve \mathcal{B}_i is the base one derived by following the order o_i and its domain is denoted D_i . In the following we establish relations between $\mathcal{V}_l(\mathcal{B}_j)$ and $\mathcal{V}_l(\mathcal{B}_i)$ so that we can bound the time complexity of step i on the entire group $G(o_i)$ (Lemma 6.11).

Before proceeding, we show that it is indeed possible to partition the block Π_i of all the $i!$ permutations of set \mathcal{S}_i in disjoint groups of i permutations each, using the scheme of (1)–(2). The proof of the following lemma was pointed out to us by Stefan Felsner in personal communication (Dec. 2019).

Lemma 6.2 *The partitioning of Π_i into disjoint groups by the scheme we defined in (1)–(2) is possible, i.e., for all $i \in \mathbb{N}$ and any block Π_i of permutations on \mathcal{S}_i there exists a set $F \subset \Pi_i$ of $(i - 1)!$ permutations such that $\Pi_i = \bigcup_{o \in F} G(o)$; that is, $G(\pi) \cap G(\sigma) = \emptyset$, for any $\pi, \sigma \in F$.*

Proof Following [15], denote by $\lfloor \pi \rfloor$ the set of all permutations that are obtained from a permutation π by deleting one element. Let $F \subseteq \Pi_i$ be a set of permutations such that $\lfloor \pi \rfloor$ and $\lfloor \sigma \rfloor$ are disjoint, for each $\pi, \sigma \in F$. Levenshtein calls such a family F of $(i - 1)!$ permutations a *code capable of correcting single deletions*, and proves that these codes exist for all $i \in \mathbb{N}$ [15, Thm. 3.1]. The set $\lfloor \pi \rfloor$ is equivalent to $G(\pi)$. Since the set F exists, it follows that Π_i is the disjoint union $\bigcup_{o \in F} G(o)$. \square

We can now proceed to estimate the time complexity of step i on one group of permutations $G(o_i)$. We first introduce some terminology.

Definition 6.3 Let α' be an auxiliary arc in \mathcal{B}_j and let $\alpha \in \mathcal{S}_i$ be a core arc of the same site. We say that α' is an auxiliary arc of α if, at step k , when $\alpha = \alpha_k$ is inserted in \mathcal{B}_j ,




Fig. 31 Illustration for Definition 6.3, $o_1 = (\beta, \alpha, \gamma)$: The core arc $\alpha \in \mathcal{S}_i$ is the source of $\alpha' \in \text{in}_1^+$. The expanded arc $\tilde{\alpha} \supseteq \alpha'$ was created at the time of inserting α , while constructing \mathcal{B}_1




Fig. 32 **a** Boundary curve \mathcal{B}_i , where $o_i = (\gamma, \beta, \alpha)$. **b** Boundary curve \mathcal{B}_1 , where $o_1 = (\beta, \alpha, \gamma)$, containing arcs $\alpha', \beta' \in \text{in}_1$, because γ was inserted last

the created original arc $\tilde{\alpha} \supseteq \alpha \cup \alpha'$ (see Fig. 31). The core arc $\alpha \in \mathcal{S}_i$ is called the *source of α'* , denoted $\text{source}_j(\alpha')$. If α' appears on $J(s, s_\alpha)$ counterclockwise (resp. clockwise) from its source α , then α' is called a *ccw* (resp. *cw*) auxiliary arc. For example in Fig. 31, α' is a *cw* auxiliary arc of α .

The source indicates the core arc in \mathcal{S}_i that creates α' . \mathcal{S}_i may contain several core arcs of the same site, but only one of them is the source of α' .

The boundary curves $\mathcal{B}_j, j < i$, may get in and out of the domain D_i , see Fig. 30. To identify their differences from \mathcal{B}_i , let $\text{in}_j = \mathcal{B}_j \cap D_i$, and $\text{out}_j = \mathcal{B}_j \setminus \overline{D_i}$, denote the portion of \mathcal{B}_j inside, and outside of D_i , respectively. We partition the auxiliary arcs of in_j into in_j^+ and in_j^- , where in_j^+ (resp. in_j^-) includes the *ccw* (resp. *cw*) auxiliary arcs of in_j , see Fig. 32. In the following we only consider in_j^+ as in_j^- is symmetric.

Observation 6.4 *The boundary curve $\mathcal{B}_j, j \neq i$, contains no auxiliary arcs of α_j , as α_j appears last in o_j . All arcs in \mathcal{B}_i appear in \mathcal{B}_j except any auxiliary arcs of α_j . No arc of out_j can have a region adjacent to $R(\alpha_j, \mathcal{B}_j)$ in $\mathcal{V}_l(\mathcal{B}_j)$.*

Proof Since the insertion order of all core arcs, except α_j , is identical in o_i and o_j , it follows that all auxiliary arcs of \mathcal{B}_i , except any auxiliary arcs of α_j , must also appear in \mathcal{B}_j .

Observe that any auxiliary arc $\alpha' \in \text{out}_j$ must lie below (as seen from D_i) an auxiliary arc in \mathcal{B}_i , by the definition of out_j , and the fact that \mathcal{B}_i and \mathcal{B}_j are defined on the same set of core arcs. Thus, α' must lie below an auxiliary arc of α_j , see Fig. 33, where α' and α'' in \mathcal{B}_1 lie below the auxiliary arcs γ' and γ'' of $\alpha_1 = \gamma$ in \mathcal{B}_i . Since arcs of the same site cannot have adjacent regions, no auxiliary arc of α_j can have a region adjacent to $R(\alpha_j, \mathcal{B}_j)$; the claim follows. \square

Observation 6.5 *Let $\alpha' \in \text{in}_j$ and let $\alpha_k = \text{source}_j(\alpha')$. Then $k > j$, i.e., α_k follows α_j in o_i . Further, if $\alpha' \in \text{in}_j^+$ then $(\alpha_k, \alpha_j, \alpha')$ appear *ccw* in \mathcal{B}_j .*




Fig. 33 **a** Boundary curve \mathcal{B}_i , where $o_i = (\gamma, \alpha, \beta, \delta)$. **b** Boundary curve \mathcal{B}_1 containing arcs α', α'' in out_1 , where $o_1 = (\alpha, \beta, \delta, \gamma)$




Fig. 34 If $\alpha', \beta' \in in_j^+$, then $j < k < \ell$ and $(\alpha_k, \alpha_\ell, \alpha_j, \beta', \alpha')$ appear in ccw order on \mathcal{B}_j

Observation 6.6 Figure 34 indicates the structure of in_j^+ . Let $\alpha', \beta' \in in_j^+$ such that $\alpha_k = source_j(\alpha')$, $\alpha_\ell = source_j(\beta')$, and $k < \ell$. Then $j < k < \ell$ and $(\alpha_k, \alpha_\ell, \alpha_j, \beta', \alpha')$ appear in ccw order along \mathcal{B}_j . Further, all auxiliary arcs of α_ℓ must appear before the auxiliary arcs of α_k as we move on \mathcal{B}_j counterclockwise from α_j .

Since many auxiliary arcs of in_j^+ can have the same source, we define

$$N_j = \{source_j(\alpha') \in \mathcal{S}_i : \alpha' \in in_j^+\}.$$

All arcs in N_j are of different sites. Sets in_j^+ and in_k^+ , $k \neq j$, may contain many common arcs, however, we have the following disjointness property.

Lemma 6.7 $N_j \cap N_k = \emptyset$, for all $k \neq j$. Thus, $\sum_{j=1}^i |N_j| = O(i)$.

Proof Suppose $\alpha_\ell \in N_j \cap N_k$ and $j < k$, then $\alpha_\ell = source_j(\alpha')$, where $\alpha' \in in_j^+$ and $\alpha_\ell = source_k(\alpha'')$, where $\alpha'' \in in_k^+$. (The arcs α' and α'' may or may not overlap). By Observation 6.5, $j < \ell$ (resp. $k < \ell$) and $(\alpha_\ell, \alpha_j, \alpha')$ (resp. $(\alpha_\ell, \alpha_k, \alpha'')$) must appear in ccw order on \mathcal{B}_j (resp. \mathcal{B}_k).

Suppose first that $(\alpha_\ell, \alpha_k, \alpha_j)$ appear in ccw order on \mathcal{B}_i . Then, since $k < \ell$, the arc α_k is inserted before α_ℓ in \mathcal{B}_j , and thus, α' cannot exist in \mathcal{B}_j , see Fig. 35. Suppose now that $(\alpha_\ell, \alpha_j, \alpha_k)$ appear in ccw order on \mathcal{B}_i . Then, since $j < \ell$, the arc α_j is inserted before α_ℓ in \mathcal{B}_k , thus, α'' cannot exist on \mathcal{B}_k , see Fig. 36. In either case we derive a contradiction. \square

Next we establish that the parameters of the time complexity analysis for step i , as given in Definition 4.6 and Lemma 4.7, sum up to $O(i)$ on all boundary curves \mathcal{B}_j , $j \leq i$.




Fig. 35 Illustration for Lemma 6.7. The case $(\alpha_\ell, \alpha_k, \alpha_j)$ appear ccw




Fig. 36 Illustration for Lemma 6.7. The case $(\alpha_\ell, \alpha_j, \alpha_k)$ appear ccw

Lemma 6.8 *Considering all the boundary curves of group $G(o_i)$,*

$$\sum_{j=1}^i (d_1(\alpha_j, \mathcal{B}_j) + d_2(\alpha_j, \mathcal{B}_j) + \tilde{d}(\alpha_j, \mathcal{B}_j)) = O(i).$$

Proof Let α and γ denote the original arcs preceding and following α_j respectively in \mathcal{B}_i (equiv. in \mathcal{B}_j). Let $d(\alpha_j, \mathcal{B}_k)$ denote the auxiliary arcs on the boundary curve \mathcal{B}_k , $k = i, j$, from α to γ .

We first observe that $d(\alpha_j, \mathcal{B}_j)$ cannot contain any portion of out_j because no auxiliary arc of α_j may appear in \mathcal{B}_i from α to γ , since α_j is the only core arc on \mathcal{B}_i between α to γ . Thus, we only need to consider the auxiliary arcs of in_j . Next, we observe that no two auxiliary arcs in $d(\alpha_j, \mathcal{B}_j)$ can have the same source in N_j for the same reason, i.e., there is no core arc from α to γ except α_j . Thus, we can bound $d(\alpha_j, \mathcal{B}_j) \leq d(\alpha_j, \mathcal{B}_i) + |N_j|$. Then, by Lemma 6.7,

$$\sum_{j=1}^i d(\alpha_j, \mathcal{B}_j) \leq |\mathcal{B}_i| + O(i) = O(i).$$

Since $d_1(\alpha_j, \mathcal{B}_j) + d_2(\alpha_j, \mathcal{B}_j) \leq d(\alpha_j, \mathcal{B}_j)$, it follows

$$\sum_{j=1}^i (d_1(\alpha_j, \mathcal{B}_j) + d_2(\alpha_j, \mathcal{B}_j)) = O(i).$$

If $\tilde{d}(\alpha_j, \mathcal{B}_j) > 0$, we have case (d) of Observation 4.1. In this case, the endpoints of α_j are incident to Γ , both in \mathcal{B}_j and \mathcal{B}_i . Then, by Observations 6.4 and 6.6, we have




Fig. 37 Illustration for Lemma 6.9. Between any two consecutive adjacencies of $R(\alpha_j, \mathcal{B}_j)$ with regions of auxiliary arcs in in_j of the same source, there must be an adjacency with an arc $\varepsilon \in \mathcal{B}_j \cap \mathcal{B}_i$

both $\text{in}_j = \emptyset$ and $\text{out}_j = \emptyset$, implying that $\mathcal{B}_j = \mathcal{B}_i$; thus, $\tilde{d}(\alpha_j, \mathcal{B}_j) = \tilde{d}(\alpha_j, \mathcal{B}_i)$. Then,

$$\sum_{j=1}^i |\tilde{d}(\alpha_j, \mathcal{B}_j)| \leq |\tilde{\mathcal{B}}_i| = O(i). \quad \square$$

Lemma 6.9 $|R(\alpha_j, \mathcal{B}_j)| \leq 2|R(\alpha_j, \mathcal{B}_i)| + |N_j|$.

Proof We compare $R(\alpha_j, \mathcal{B}_j)$ and $R(\alpha_j, \mathcal{B}_i)$ and bound differences in their adjacencies. By Observation 6.4 no arc in out_j can have a region adjacent to $R(\alpha_j, \mathcal{B}_j)$. We also observe the following: if an arc $\varepsilon \in \mathcal{B}_j \cap \mathcal{B}_i$, common to both \mathcal{B}_j and \mathcal{B}_i , has a region $R(\varepsilon, \mathcal{B}_j)$ adjacent to $R(\alpha_j, \mathcal{B}_j)$ in $\mathcal{V}_l(\mathcal{B}_j)$, then $R(\varepsilon, \mathcal{B}_i)$ must also be adjacent to $R(\alpha_j, \mathcal{B}_i)$ in $\mathcal{V}_l(\mathcal{B}_i)$, see Fig. 37. This is correct, because otherwise, the Voronoi edge e bounding $R(\alpha_j, \mathcal{B}_j)$ and $R(\varepsilon, \mathcal{B}_j)$ (or a portion of it) would be contained in a region $R(\eta, \mathcal{B}_i)$ for an arc η that does not appear in \mathcal{B}_j , i.e., $\eta \in \text{out}_j$. By Observation 6.4, this arc may only be an auxiliary arc of α_j . However, by Lemma 5.1, if we insert η to $\mathcal{V}_l(\mathcal{B}_j)$, the region $R(\eta, \mathcal{B}_j \oplus \eta)$ will contain a portion of the edge e , thus, it will be adjacent to $R(\alpha_j, \mathcal{B}_j \oplus \eta)$, deriving a contradiction, as arcs of the same site cannot have adjacent regions.

Let $|R(\alpha_j, \mathcal{B}_j)|_x$ denote the number of additional adjacencies that $R(\alpha_j, \mathcal{B}_j)$ may have over $R(\alpha_j, \mathcal{B}_i)$, i.e., $|R(\alpha_j, \mathcal{B}_j)| \leq |R(\alpha_j, \mathcal{B}_i)| + |R(\alpha_j, \mathcal{B}_j)|_x$. We show that $|R(\alpha_j, \mathcal{B}_j)|_x \leq |R(\alpha_j, \mathcal{B}_i)| + |N_j|$. Since auxiliary arcs of the same site can never have adjacent regions, it follows that between any two possible new adjacencies of $R(\alpha_j, \mathcal{B}_j)$ with auxiliary arcs of the same source in in_j , there must be an adjacency with some arc that is common to both \mathcal{B}_i and \mathcal{B}_j .

Since by Observation 6.6 auxiliary arcs of one source in N_j must appear in a certain order along \mathcal{B}_j , and they cannot alternate, the bound follows. \square

Lemma 6.10 Consider case (c) of Observation 4.1 at the insertion of α_j in \mathcal{B}_j . Suppose that the insertion of α_j splits an existing arc ω into two pieces ω_1 and ω_2 . Then at least one of these two arcs (say ω_1) must also exist in \mathcal{B}_i . Further, $|R(\omega_1, \mathcal{B}_j)| \leq 2|R(\omega_1, \mathcal{B}_i)| + |N_j|$.




Fig. 38 Illustration for the proof of Lemma 6.10. If $\omega_2 \notin \mathcal{B}_i$, then $\omega_1 \in \mathcal{B}_i$

Proof Suppose $\omega_1\alpha_j\omega_2$ appear in \mathcal{B}_j in ccw order and $\omega_2 \notin \mathcal{B}_i$. Then $\omega_2 \in \text{in}_j^+$, see Fig. 38. Let $\alpha_\ell = \text{source}_j(\omega_2)$, then $\ell > j$ as $\omega_2 \in \text{in}_j^+$. We claim that ω_1 must belong to \mathcal{B}_i .

Let $\tilde{\omega} \supset \alpha_\ell$ denote the expanded arc created at the insertion time of α_ℓ following the order o_j . Clearly, $\tilde{\omega} \supset \omega$. Let $\hat{\omega} \supset \alpha_\ell$ denote the expanded arc created at the insertion time of α_ℓ , following o_i . Since $\ell > j$, it follows that $\hat{\omega}$ can extend ccw at most until α_j and $\hat{\omega} \subset \tilde{\omega}$. Since $\tilde{\omega}$ extends ccw past α_j , it follows that no core arc α_ρ , with $\rho < \ell$ can exist between α_i and α_j . Thus, $\hat{\omega}$ must extend ccw to α_j and $\hat{\omega} \supset \omega_1$. In addition, no α_ρ , with $\rho > \ell$, can delete ω_1 during its insertion, while following o_i , because the same would happen in o_j and ω_1 exists in \mathcal{B}_j . Thus, ω_1 must exist in \mathcal{B}_i .

We can now bound $|R(\omega_1, \mathcal{B}_j)| \leq 2|R(\omega_1, \mathcal{B}_i)| + |N_j|$ analogously to Lemma 6.9. The only additional argument needed for the fact that no arc in out_j can have a region adjacent to $R(\omega_1, \mathcal{B}_j)$ is the observation that each arc in out_j lies below the s_ω -bisector, because arc α_j splits arc ω (case (c) of Observation 4.1). \square

Let $T(i, o_j)$ denote the time complexity of step- i following permutation o_j , i.e., the time required by the last arc insertion of o_j .

Lemma 6.11 *The time for step- i on the entire group $G = G(o_i)$ is*

$$T(i, G) = \sum_{o_j \in G} T(i, o_j) = O(i).$$

Proof Lemmas 6.9 and 6.10 establish that $|R(\alpha_j, \mathcal{B}_j)| + |R(\omega_j, \mathcal{B}_j)| \leq 2(|R(\alpha_j, \mathcal{B}_i)| + |R(\omega_j, \mathcal{B}_i)| + |N_j|)$, where ω_j denotes one of the two arcs that is split and belongs to \mathcal{B}_i , if case (c) of Observation 4.1 is concerned. Since ω_j is always an immediate neighbor of α_j , we count it at most twice, and thus, the total complexity $\sum_{j=1}^i |R(\omega_j, \mathcal{B}_i)|$ is $O(i)$. Together with Lemma 6.7 this directly implies that $\sum_{j=1}^i (|R(\alpha_j, \mathcal{B}_j)| + r(\alpha_j, \mathcal{B}_j)) = O(i)$. Lemma 6.8 establishes that

$$\sum_{j=1}^i (d_1(\alpha_j, \mathcal{B}_j) + d_2(\alpha_j, \mathcal{B}_j) + \tilde{d}(\alpha_j, \mathcal{B}_j)) = O(i).$$

Then by Lemma 4.7 the claim is derived. \square

All permutations at level- i of the decision tree are equally likely. By Lemma 6.2, it is possible to partition them into groups of i nodes each, which satisfy our scheme

of (1)–(2). By Lemma 6.11, each group requires total $O(i)$ time to perform step i on all its permutations. We thus conclude:

Theorem 6.12 *The time complexity of step i of the randomized algorithm is expected $O(1)$.*

We conclude with the following theorem.

Theorem 6.13 *Given an abstract Voronoi diagram $\mathcal{V}(S)$, the diagram $\mathcal{V}(S \setminus \{s\}) \cap \text{VR}(s, S)$ can be computed in expected $O(h)$ time, where h is the complexity of $\partial \text{VR}(s, S)$. Thus, the updated Voronoi diagram $\mathcal{V}(S \setminus \{s\})$ can be computed from $\mathcal{V}(S)$, after the deletion of site s , in expected linear time $O(h)$.*

7 Computing the Order- k Voronoi Diagram Iteratively

Our algorithm to perform deletion in expected linear-time can be adapted to iteratively compute the order- k abstract Voronoi diagram, for increasing values of k , in total time $O(k(n - k)n + n \log n)$, if $k \leq n/2$. In particular, given a face f of an order- k Voronoi region, we can compute the order- $(k + 1)$ subdivision within f in expected time $O(|\partial f|)$. In this section we describe the required adaptation over site-deletion.

The *order- k abstract Voronoi region* of a subset of sites $H \subset S$, $|H| = k$, is defined [3] as

$$\text{VR}_k(H, S) = \bigcap_{\substack{q \in H \\ p \in S \setminus H}} D(q, p).$$

The *order- k abstract Voronoi diagram* of S is [3]

$$\mathcal{V}_k(S) = \mathbb{R}^2 \setminus \bigcup_{\substack{H \subset S \\ |H|=k}} \text{VR}_k(H, S).$$

The combinatorial complexity of $\mathcal{V}_k(S)$ is $O(k(n - k))$. For $k = 1$, it is the nearest-neighbor abstract Voronoi diagram $\mathcal{V}(S)$, and for $k = n - 1$, it is the farthest abstract Voronoi diagram $\text{FVD}(S)$. The vertices of the diagram are classified into *new* and *old*, where a *new* vertex in $\mathcal{V}_k(S)$ is an *old* vertex of $\mathcal{V}_{k+1}(S)$.

Consider a face f of an order- k Voronoi region $\text{VR}_k(H)$, $H \subset S$, $|H| = k$. Let $S_f \subseteq S \setminus H$ denote the set of sites, which together with H , induce the Voronoi edges on the boundary ∂f . Our goal is to compute the Voronoi diagram of $S \setminus H$ within f , $\mathcal{V}(S_f) \cap f$, in expected linear time, i.e., in time $O(|\partial f|)$. This diagram is a tree (or forest if f is unbounded) with properties analogous to Lemma 2.1 (see also [5]). To extend Theorem 6.13 from $k = 1$ to an arbitrary k , there is a non-trivial challenge to overcome: the complexity of the boundary ∂f depends not only on $|S_f|$ but also on k . Thus, a direct application of our deletion algorithm would not result in a linear-time scheme, if k is not a constant.

Consider a face f of $\text{VR}_k(H, S)$ and its boundary ∂f . We call any piece of ∂f between two consecutive *new* vertices, an *order- k arc*. Such an arc does not have constant complexity but may contain a sequence of old Voronoi vertices on ∂f . In this section, let \mathcal{S} denote the collection of the order- k arcs along the boundary of f .

An order- k arc α , bounding the face f , is a piece of the so-called *Hausdorff bisector* between site $s_\alpha \in S_f$ and set H (see, e.g., [18] for the definition of the concrete Hausdorff bisector between two point sets). In abstract terms, the *Hausdorff bisector* between s_α and H is the boundary of the farthest Voronoi region $\text{FVR}(s_\alpha, H \cup \{s_\alpha\})$, where $\text{FVR}(s, S') = \bigcap_{q \in S' \setminus \{s\}} D(q, s)$.

Let the *Hausdorff bisector* between a site $s_\alpha \in S_f$ and H , which is relevant to face f , be defined as

$$J(s_\alpha, H) = \partial \text{FVR}(\alpha, H \cup \{s_\alpha\}),$$

where $\text{FVR}(\alpha, H \cup \{s_\alpha\})$ denotes the face of region $\text{FVR}(s_\alpha, H \cup \{s_\alpha\})$ that is incident to arc α . $J(s_\alpha, H)$ is an unbounded Jordan curve dividing the plane in two open parts; let $D(s_\alpha, H) = \text{FVR}(\alpha, H \cup \{s_\alpha\})$.

The complexity of $J(s_\alpha, H)$ is $\Theta(|H|)$, and this is an obstacle to our randomized linear time scheme. It is possible to overcome this problem by considering relaxed Hausdorff bisectors whose complexity depends solely on order- k arcs, and which define a series of even larger shrinking domains enclosing the face f . Let $H_\alpha \subseteq H$ be the subset of sites in H that, together with s_α , define the edges and vertices along the arc α . Instead of $J(s_\alpha, H)$, which is hard to compute, we consider the Hausdorff bisector $J(s_\alpha, H_\alpha)$, where $\alpha \subseteq J(s_\alpha, H_\alpha)$, and has complexity $\Theta(|H_\alpha|)$. In fact, $\alpha \subseteq J(s_\alpha, \tilde{H}_\alpha)$, for any $H_\alpha \subseteq \tilde{H}_\alpha \subseteq H$. Let $|\alpha|$ denote the complexity of arc α , $|\alpha| = |H_\alpha|$. We make use of the following property.

Lemma 7.1 $J(s_\alpha, H) \subseteq \overline{D(s_\alpha, \tilde{H}_\alpha)} \subseteq \overline{D(s_\alpha, H_\alpha)}$, where $H_\alpha \subseteq \tilde{H}_\alpha \subseteq H$.

Proof Since $H_\alpha \subseteq H$, we have

$$D(s_\alpha, H) = \text{FVR}(s_\alpha, H \cup \{s_\alpha\}) \subseteq \text{FVR}(s_\alpha, H_\alpha \cup \{s_\alpha\}) = D(s_\alpha, H_\alpha). \tag{3}$$

Thus, it holds $J(s_\alpha, H) = \partial D(s_\alpha, H) \subseteq \overline{D(s_\alpha, H_\alpha)}$. Analogously we can show the subset relation for \tilde{H}_α . □

It is now straightforward to adapt the algorithm of Sect. 6, using appropriate Hausdorff bisectors that are derived by the order- k arcs in \mathcal{S} , in place of the s -related bisectors in the previous sections. The complexity of each such Hausdorff bisector must be proportional to the complexity of its underlying order- k arc. Lemma 7.1 implies the correctness of this adaptation.

We start with domain D_1 defined by $J(s_{\alpha_1}, H_{\alpha_1})$, i.e., $D_1 = D(s_{\alpha_1}, H_{\alpha_1}) \cap D_\Gamma$, for the first order- k arc α_1 of a random permutation of \mathcal{S} . The boundary complexity of D_1 is $O(|\alpha_1|)$.

Note that D_1 is a superset of domain $D(s_{\alpha_1}, H) \cap D_\Gamma$. At step i , we insert arc α_i considering bisector $J(s_{\alpha_i}, \tilde{H}_{\alpha_i})$, where $H \supseteq \tilde{H}_{\alpha_i} \supseteq H_{\alpha_i}$, and $|\tilde{H}_{\alpha_i}| \leq |H_{\alpha_i}| + 2$. We

use \tilde{H}_{α_i} , possibly a superset of H_{α_i} , in order to include at most one site in H for each neighbor of α_i in \mathcal{P}_i . This is done to correctly link two neighboring order- k arcs on \mathcal{P}_i so that they are both incident to a common (new) Voronoi vertex. By Lemma 7.1, domain D_i is a superset of the domain we would get if we instead considered bisector $J(s_{\alpha_i}, H) \supset \alpha_i$. Therefore, the relaxed construction works correctly. At the end, $D_h = f$.

We conclude that Theorem 6.13 applies, constructing $\mathcal{V}(\mathcal{S}) = \mathcal{V}(S_f) \cap f$ in expected time $O(|\partial f|)$.

Since the complexity of $\mathcal{V}_k(S)$ is $O(k(n - k))$, the $O(k^2(n - k) + n \log n)$ bound for iteratively constructing the diagram, starting at $\mathcal{V}(S)$, easily follows for $k \leq n/2$. Although there are algorithms of better time complexity to construct $\mathcal{V}_k(S)$, such as the $O(k(n - k) \log^2 n + n \log^3 n)$ randomized incremental algorithm of Bohler et al. [5], the iterative construction is nice and simple, therefore, it can be preferable for small values of k .

8 The Farthest Abstract Voronoi Diagram

In this section we show how to modify (in fact simplify) the algorithm for the deletion of one site to compute the *farthest* abstract Voronoi diagram, after the sequence of its faces at infinity is known.

The *farthest Voronoi region* of a site $p \in S$ is $FVR(p, S) = \bigcap_{q \in S \setminus \{p\}} D(q, p)$ and the *farthest abstract Voronoi diagram* of S is $FVD(S) = \mathbb{R}^2 \setminus \bigcup_{p \in S} FVR(p, S)$. $FVD(S)$ is a tree of complexity $O(n)$, however, regions may be disconnected and a farthest Voronoi region may consist of $\Theta(n)$ disjoint faces [16]. Let $D^*(p, q) = D(q, p)$; then $FVR(p, S) = \bigcap_{q \in S \setminus \{p\}} D^*(p, q)$.

Unless otherwise noted, we adopt the following convention: we reverse the labels of bisectors and use $D^*(\cdot, \cdot)$, in the place of $D(\cdot, \cdot)$, in most definitions and constructs of Sects. 3 and 4. Under this convention the definition of e.g., a p -monotone path remains the same but it uses $\partial FVR(p, \cdot)$ in the place of $\partial VR(p, \cdot)$. The corresponding arrangement of p -related bisectors $\mathcal{J}_{p, S'}$, $S' \subseteq S$, is considered with the labels of bisectors and their dominance regions reversed from the original system \mathcal{J} .

Consider the enclosing curve Γ as defined in Sect. 2, and let \mathcal{S} be the sequence of arcs on Γ derived by $\Gamma \cap FVD(S)$. \mathcal{S} represents the sequence of the farthest Voronoi faces in $FVD(S)$ at infinity. The domain of computation is D_Γ . For an arc α of \mathcal{S} , let s_α denote the site in S for which $\alpha \subset FVR(s_\alpha, S)$. With respect to site occurrences, \mathcal{S} is a Davenport–Schinzel sequence of order 2. \mathcal{S} can be computed in time $O(n \log n)$ in a divide and conquer fashion, similarly to computing the *hull* of a farthest segment Voronoi diagram, see e.g., [19].

We treat the arcs in \mathcal{S} as sites and compute $\mathcal{V}(\mathcal{S}) = FVD(S) \cap D_\Gamma$. Let $VR(\alpha, \mathcal{S})$ denote the face of $FVD(S) \cap D_\Gamma$ incident to $\alpha \in \mathcal{S}$, see Fig. 39. $\mathcal{V}(\mathcal{S})$ is a tree whose leaves are the endpoints of the arcs in \mathcal{S} .

Consider $S' \subseteq \mathcal{S}$, and let $S' \subseteq S$ be the set of sites that define the arcs in S' .

Definition 8.1 A *boundary curve* \mathcal{P} for S' is a partitioning of Γ into *arcs* by the bisector system $\mathcal{J}_{s, S'}$, such that any two consecutive arcs $\alpha, \beta \in \mathcal{P}$ are incident to $J(s_\alpha, s_\beta) \in \mathcal{J}_{s, S'}$, having *consistent labels*, and \mathcal{P} contains an arc $\alpha \supseteq \alpha^*$, for every




Fig. 39 The farthest Voronoi diagram $\mathcal{V}(S) = \text{FVD}(S) \cap D_\Gamma$ and the Voronoi region $\text{VR}(\alpha, S)$. Bisector labels are shown in the farthest (reversed) sense

core arc $\alpha^* \in S'$. We say that the labels of α, β are consistent, if there is a neighborhood $\tilde{\alpha} \subseteq \alpha$ and $\tilde{\beta} \subseteq \beta$ incident to the common endpoint of α and β such that $\tilde{\alpha} \in D^*(s_\alpha, s_\beta)$ and $\tilde{\beta} \in D^*(s_\beta, s_\alpha)$.

There can be several different boundary curves for S' . The arcs in \mathcal{P} that contain a core arc in S' are called *original* and any remaining arcs are called *auxiliary*. The arcs in \mathcal{P} , although they are arcs on Γ , they are all boundary arcs and none is considered a Γ -arc in the sense of the previous sections. The endpoint $J(s_\alpha, s_\beta) \cap \Gamma$ on \mathcal{P} separating two consecutive arcs α, β is denoted by $\nu(\alpha, \beta)$.

The Voronoi-like diagram of a boundary curve \mathcal{P} is defined analogously to Definition 3.3. Since \mathcal{P} consists only of boundary arcs, $\mathcal{V}_l(\mathcal{P})$ is a tree whose leaves are the vertices of \mathcal{P} . The properties of a Voronoi-like diagram in Sect. 3 remain the same (under the conventions of this section).

Given $\mathcal{V}_l(\mathcal{P})$ for a boundary curve \mathcal{P} of $S' \subset S$, we can insert a core arc $\beta^* \in S \setminus S'$ and obtain $\mathcal{V}_l(\mathcal{P} \oplus \beta^*)$. The insertion is performed analogously to Sect. 4. The original arc $\beta \supseteq \beta^*$, with endpoints x, y is defined as follows: let δ be the first arc on \mathcal{P} counterclockwise (resp. clockwise) from β^* such that $J(s_\beta, s_\delta) \cap \delta \neq \emptyset$; let $x = \nu(\delta, \beta)$ (resp. $y = \nu(\beta, \delta)$). Let $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$ be the boundary curve obtained from \mathcal{P} by substituting with β its overlapping piece from x to y . No original arc of \mathcal{P} can be deleted by the insertion of β . Observation 4.1 remains the same, except from cases (d) and (e) which do not exist.

The *merge curve* $J(\beta)$, given $\mathcal{V}_l(\mathcal{P})$, is defined analogously to Definition 4.2; it is only simpler as it does not contain Γ -arcs. Theorem 4.3 remains valid, i.e., $J(\beta)$ is an s_β -monotone path in $\mathcal{J}_{s_\beta, S'}$ connecting the endpoints of β . The proof structure is the same as for Theorem 4.3, however, Lemma 4.11 now requires a different proof, which we give in the sequel (see Lemma 8.3). Lemma 4.12 is not relevant; while Lemmas 4.13 and 4.14 are analogous.

In the following lemma we restore the labeling of bisectors to the original.

Lemma 8.2 *In an admissible bisector system \mathcal{J} , there cannot be two p -cycles, $p \in S$, with disjoint interior.*

Proof By its definition, the nearest Voronoi region $\text{VR}(p, S)$ (resp. $\text{VR}(p, S) \cap D_\Gamma$) must be enclosed in the interior of any p -cycle of the admissible bisector system \mathcal{J}




Fig. 40 Illustration for Lemma 8.3. Nearest labels are shown

(resp. $\mathcal{J} \cup \{\Gamma\}$). But $\text{VR}(p, S)$ (resp. $\text{VR}(p, S) \cap D_\Gamma$) is connected (by axiom (A1)), thus, there cannot be two different p -cycles with disjoint interior. \square

Lemma 8.3 Consider the merge curve $J(\beta)$. Suppose v_{i+1} is not a valid vertex because $v_{i+1} \in \alpha_i$, i.e., e_i hits arc α_i . Then vertex v_{m-j} cannot be on \mathcal{P} .

Proof Suppose otherwise, i.e., vertex v_{m-j} is on the boundary arc α_{m-j} . Then J_x^i and J_y^j partition D_Γ in three parts: a middle part incident to β , and two parts C_1 and C_2 at either side of J_x^i and J_y^j respectively, whose closures are disjoint, see Fig. 40. But the boundaries of C_1 and C_2 are s_β -cycles in the admissible bisector system $\mathcal{J} \cup \{\Gamma\}$ contradicting Lemma 8.2. Note that here we use the original labels of bisectors, including $\Gamma = J(s_\beta, s_\infty)$. \square

The diagram $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ is defined analogously and the proof that $\mathcal{V}_l(\mathcal{P}) \oplus \beta$ is the Voronoi-like diagram $\mathcal{V}_l(\mathcal{P}_\beta)$ for $\mathcal{P}_\beta = \mathcal{P} \oplus \beta$, is analogous to the proof of Theorem 4.4.

The randomized algorithm for computing $\mathcal{V}(S) = \text{FVD}(S) \cap D_\Gamma$ is the same as in Sect. 6. The time analysis is also completely analogous. For completeness we point out that, here, the set out_j consists of the auxiliary arcs in \mathcal{B}_j that overlap with the auxiliary arcs of α_j in \mathcal{B}_i . The set in_j are any remaining auxiliary arcs in $\mathcal{B}_j \setminus \text{out}_j$ that differ from the corresponding auxiliary arcs in \mathcal{B}_i . All observations of Sect. 6.1 remain intact under this updated notion of in_j and out_j . Thus, the (expected) linear time complexity can be analogously established.

Theorem 8.4 Given the sequence of its faces at infinity, i.e., given the sequence of arcs S implied by $\text{FVD}(S) \cap \Gamma$, the farthest abstract Voronoi diagram $\text{FVD}(S)$ can be computed in expected linear time $O(|S|)$.

9 Concluding Remarks

In this paper we formalized the notion of an *abstract Voronoi-like diagram*, which is defined as a tree (or forest) on the arrangement of the underlying bisector system related to a set of abstract sites S . We defined the Voronoi-like diagram of a *boundary curve*, which is implied by a subset of Voronoi edges bounding a Voronoi region $\text{VR}(s, S)$. We showed that the Voronoi-like diagram of a boundary curve is well

defined, unique, and robust under an arc-insertion operation, which enables its use in incremental constructions. Using Voronoi-like diagrams as intermediate structures, we derived a very simple, randomized incremental algorithm to update an abstract Voronoi diagram, after deletion of one site, in expected linear time. The algorithm is applicable to any concrete diagram that falls under the umbrella of abstract Voronoi diagrams. In addition, the time complexity analysis offers a variant to backwards analysis, applicable to order-dependent structures.

The technique can be adapted to compute the order- $(k+1)$ subdivision within an order- k abstract Voronoi region, and the farthest abstract Voronoi diagram, after the order of its faces at infinity is known. The Voronoi-like structure provides the means to deal with the underlying disconnected Voronoi regions, which is the common complication of these, otherwise simple, Voronoi structures.

A deterministic linear-time construction of these diagrams remains an open problem. In future work we would like to consider Voronoi-like structures in the linear-time framework of Aggarwal et al. [1] aiming at a deterministic linear-time algorithm for the same problems.

Acknowledgements We sincerely thank Stefan Felsner for the proof of Lemma 6.2 making the connection to the seemingly unrelated result of Levenshtein [15] on perfect codes, which established this claim for the time complexity analysis.

Funding Open access funding provided by Università della Svizzera italiana.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aggarwal, A., Guibas, L.J., Saxe, J., Shor, P.W.: A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete Comput. Geom.* **4**(6), 591–604 (1989)
2. Aurenhammer, F., Klein, R., Lee, D.-T.: *Voronoi Diagrams and Delaunay Triangulations*. World Scientific, Hackensack (2013)
3. Bohler, C., Cheilaris, P., Klein, R., Liu, Ch.-H., Papadopoulou, E., Zavershynskiy, M.: On the complexity of higher order abstract Voronoi diagrams. *Comput. Geom.* **48**(8), 539–551 (2015)
4. Bohler, C., Klein, R., Lingas, A., Liu, Ch.-H.: Forest-like abstract Voronoi diagrams in linear time. *Comput. Geom.* **68**, 134–145 (2018)
5. Bohler, C., Klein, R., Liu, Ch.-H.: An efficient randomized algorithm for higher-order abstract Voronoi diagrams. *Algorithmica* **81**(6), 2317–2345 (2019)
6. Buchin, K., Devillers, O., Mulzer, W., Schrijvers, O., Shewchuk, J.: Vertex deletion for 3D Delaunay triangulations. In: 21st Annual European Symposium on Algorithms (Sophia Antipolis 2013). *Lecture Notes in Comput. Sci.*, vol. 8125, pp. 253–264. Springer, Heidelberg (2013)
7. Chew, P.L.: Building Voronoi diagrams for convex polygons in linear expected time. Technical report PCS-TR90-147, Dartmouth College (1990). https://digitalcommons.dartmouth.edu/cs_tr/47/
8. Chin, F., Snoeyink, J., Wang, C.A.: Finding the medial axis of a simple polygon in linear time. *Discrete Comput. Geom.* **21**(3), 405–420 (1999)

9. Junginger, K., Papadopoulou, E.: Deletion in abstract Voronoi diagrams in expected linear time. In: 34th International Symposium on Computational Geometry (Budapest 2018). Leibniz Int. Proc. Inform., vol. 99, # 50. Leibniz-Zent. Inform., Wadern (2018)
10. Khramtcova, E., Papadopoulou, E.: An expected linear-time algorithm for the farthest-segment Voronoi diagram (2017). [arXiv:1411.2816v3](https://arxiv.org/abs/1411.2816v3)
11. Klein, R.: Concrete and Abstract Voronoi Diagrams. Lecture Notes in Comput. Sci., vol. 400. Springer, Berlin (1989)
12. Klein, R., Langetepe, E., Nilforoushan, Z.: Abstract Voronoi diagrams revisited. *Comput. Geom.* **42**(9), 885–902 (2009)
13. Klein, R., Lingas, A.: Hamiltonian abstract Voronoi diagrams in linear time. In: 5th International Symposium on Algorithms and Computation (Beijing 1994). Lecture Notes in Comput. Sci., vol. 834, pp. 11–19. Springer, Berlin (1994)
14. Klein, R., Mehlhorn, K., Meiser, S.: Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom.* **3**(3), 157–184 (1993)
15. Levenshtein, V.I.: Perfect codes in the metric of deletions and insertions. *Discrete Math. Appl.* **2**(3), 241–258 (1992)
16. Mehlhorn, K., Meiser, S., Rasch, R.: Furthest site abstract Voronoi diagrams. *Int. J. Comput. Geom. Appl.* **11**(6), 583–616 (2001)
17. Okabe, A., Boots, B., Sugihara, K., Chiu, S.N.: *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley Series in Probability and Statistics. Wiley, Chichester (2000)
18. Papadopoulou, E.: The Hausdorff Voronoi diagram of point clusters in the plane. *Algorithmica* **40**(2), 63–82 (2004)
19. Papadopoulou, E., Dey, S.K.: On the farthest line-segment Voronoi diagram. *Int. J. Comput. Geom. Appl.* **23**(6), 443–459 (2013)
20. Seidel, R.: Backwards analysis of randomized geometric algorithms. In: *New Trends in Discrete and Computational Geometry*. Algorithms Combin., vol. 10, pp. 37–67. Springer, Berlin (1993)
21. Sharir, M., Agarwal, P.K.: *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, Cambridge (1995)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.