Stabbing circles for sets of segments in the plane*

Mercè Claverol[1], Elena Khramtcova[2], Evanthia Papadopoulou[2], Maria Saumell[3], and Carlos Seara[1]

[1] Universitat Politècnica de Catalunya, Spain.
[2] Faculty of Informatics, Università della Svizzera italiana (USI), Lugano, Switzerland.
[3] Department of Mathematics and European Centre of Excellence NTIS, University of West Bohemia, Czech Republic.

Abstract. Stabbing a set S of n segments in the plane by a line is a well-known problem. In this paper we consider the variation where the stabbing object is a circle instead of a line. We show that the problem is tightly connected to two cluster Voronoi diagrams, in particular, the Hausdorff and the farthest-color Voronoi diagram. Based on these diagrams, we provide a method to compute a representation of all the combinatorially different stabbing circles for S, and the stabbing circles with maximum and minimum radius. We give conditions under which our method is fast. These conditions are satisfied if the segments in S are parallel, resulting in a $O(n \log^2 n)$ time and $O(n)$ space algorithm. We also observe that the stabbing circle problem for S can be solved in worst-case optimal $O(n^2)$ time and space by reducing the problem to computing the stabbing planes for a set of segments in 3D. Finally we show that the problem of computing the stabbing circle of minimum radius for a set of n parallel segments of equal length has an $\Omega(n \log n)$ lower bound.

1 Introduction

Let S be a set of n line segments (segments for short) in the plane. We say that a region $\mathcal{R} \subseteq \mathbb{R}^2$ is a *stabbing region* for S if *exactly* one endpoint of each segment of S lies in the exterior of $\mathcal{R}$. The boundary of $\mathcal{R}$ (also known as a *stabber* for S) intersects all the segments in S and separates/classifies their endpoints into two classes, depending on whether or not they lie in the exterior of $\mathcal{R}$. Two stabbing regions $\mathcal{R}_1$ and $\mathcal{R}_2$ for S are *combinatorially different* if they classify the endpoints of S differently.

A natural problem is to determine the existence and compute (when possible) a representation of all combinatorially different stabbing regions for S. We are interested in stabbing regions whose boundary has constant complexity. Perhaps the simplest such region is a halfplane bounded by a line that intersects or *stabs* all the segments. Edelsbrunner et al. [15] presented an optimal $\Theta(n \log n)$ time algorithm to compute a representation of all the $O(n)$ combinatorially different stabbing lines for S. An $\Omega(n \log n)$ lower bound for the decision problem was later presented by Avis et al. [6]. For parallel segments the problem can be solved in $O(n)$ time by linear programming. If no stabbing halfplane for S exists, it is natural to ask for other types of stabbers. Computing

* A preliminary version of this paper appeared in Proc. 12th Latin American Theoretical Informatics Symposium (LATIN'16), pp. 290–305.

all the combinatorially different stabbing wedges (regions defined by the intersection of two halfplanes) can be carried out in $O(n^3 \log n)$ time and $O(n^2)$ space [11]. The same question can be answered in $O(n \log n)$ time and $O(n)$ space for isothetic stabbing strips, quadrants and 3-sided rectangles; and in $O(n^2 \log n)$ time and $O(n^2)$ space for isothetic stabbing rectangles [12].

In this paper, we focus on the *stabbing circle problem*, formulated as follows. Let S be a set of n segments in the plane in general position (segments have $2n$ distinct endpoints, no three endpoints are collinear, and no four of them are cocircular). A circle c is a stabbing circle for S if exactly one endpoint of each segment in S is contained in the exterior of the closed disk (region) induced by c; see Figure 1. The stabbing circle problem for S consists of (1) answering whether a stabbing circle for S exists; (2) reporting a representation (for the centers) of all the combinatorially different stabbing circles for S; and (3) finding stabbing circles with minimum and maximum radius. Note that stabbing circles of minimum radius do not always exist. On the contrary, there are cases in which any stabbing circle can be shrunk by decreasing its radius or moving its center, however, the "limit" circle is not stabbing anymore. In such cases, our task is to find this "limit" circle. The same may happen with the stabbing circles of maximum radius; refer to Lemma 21 for the details. We remark that this can happen even if the general position assumptions used in this paper are fulfilled. Note also that our stabbing criterion uses only the segment endpoints, thus, S can be seen as a set of pairs of points, where a segment is simply a convenient representation for such a pair.

Fig. 1. Left: Segment set with a stabbing circle. Right: Segment set with no stabbing circle.

Other works with similar criteria are as follows: Rappaport [25] considers the problem of computing the stabbing simple polygon of minimum perimeter for a set S of general segments, where a simple polygon stabs S if *at least* one point (which is not necessarily an endpoint) of each segment is in the polygon; this minimum stabbing polygon is always a convex polygon. Díaz-Báñez et al. [14] focus on computing the stabbing simple polygons of minimum perimeter or area with a distinct criterion, specifically, that at least one endpoint of each segment is required to be in the polygon. Keeping the same criterion and replacing simple polygons by disks leads to the so-called *pairs of points L_2 1-center problem*, studied by Arkin et al. [3]; the goal of the problem is to determine a minimum-radius disk that contains at least one endpoint of each segment. Finally, Arkin et al. [2] consider, given a collection of compact sets, whether there exists a *convex* body $\mathcal{R}$ whose boundary intersects every set in the collection. They show that, for segment sets, deciding the existence of a convex stabber is NP-hard.

Our results. First, we point out a connection between the stabbing circle problem and two *cluster Voronoi diagrams*: the *Hausdorff* and the *farthest-color Voronoi diagram*. This connection is interesting in its own right and it forms the base of our method to solve the stabbing circle problem. For a family of clusters (sets) of points, the Hausdorff Voronoi diagram (HVD) is a subdivision of the plane into regions such that every point within one region has the same nearest cluster, where the distance between a point $p \in \mathbb{R}^2$ and a cluster C is the maximum distance between p and all points in C. The farthest-color Voronoi diagram (FCVD) is the reverse: it reveals the farthest cluster for every point in a region, according to the minimum distance between a point and a cluster. Both diagrams have quadratic structural complexity in the worst case [1,16,19]. However, for some classes of input sites the diagrams are linear and can be constructed efficiently, see e.g. [8,23] for the HVD. Here, clusters are the pairs of segment endpoints, and S is a family of such pairs of points.

Our central object is $\mathsf{FCVD}^{\star}(S)$, defined as the locus of points whose farthest-color neighbor (i.e., their *owner* in the farthest-color Voronoi diagram) is closer than their nearest cluster (i.e., their *owner* in the Hausdorff Voronoi diagram). We observe that any point $p \in \mathbb{R}^2$ that is the center of a stabbing circle for S lies in the interior of $\mathsf{FCVD}^{\star}(S)$. The points in the interior of $\mathsf{FCVD}^{\star}(S)$ that are not centers of stabbing circles for S are separators among centers of combinatorially different stabbing circles.

Thus, $\mathsf{FCVD}^{\star}(S)$ provides all the information that is relevant to stabbing circles: whether such circles exist, a list of all combinatorially different stabbing circles, and the stabbing circles with minimum and maximum radius. We identify sufficient conditions for efficient algorithms to construct $\mathsf{FCVD}^{\star}(S)$, and thus, to solve the stabbing circle problem. These conditions are: (1) the Hausdorff Voronoi diagram and the farthest-color Voronoi diagram have linear structural complexity and can be constructed fast; (2) the edges of the Hausdorff Voronoi diagram are not "spoiled" many times, where by "spoiling" an edge e we mean a technical condition necessary to cause $e \cap \mathsf{FCVD}^{\star}(S)$ to be disconnected. If the segments in S are parallel, conditions (1) and (2) are satisfied, and we obtain that the stabbing circle problem for S can be solved in $O(n \log^2 n)$ time and $O(n)$ space. As a byproduct, we establish that the farthest-color Voronoi diagram for such a set S has structural complexity $O(n)$ and can be constructed in $O(n \log n)$ time and $O(n)$ space, which was not previously known. In addition, we show that the problem of computing the stabbing circle of minimum radius for a set of n parallel segments of equal length has an $\Omega(n \log n)$ lower bound.

Summary. In Section 2 we give the necessary definitions; in addition, we observe that, using a known technique, the stabbing circle problem for arbitrary segments can be solved in $O(n^2)$ time and space. In Section 3 we show the connection of $\mathsf{FCVD}^{\star}(S)$ with the problem, and we give useful properties of $\mathsf{HVD}(S)$, $\mathsf{FCVD}(S)$, and $\mathsf{FCVD}^{\star}(S)$. In Section 4 we present an algorithm to compute $\mathsf{FCVD}^{\star}(S)$. In Section 5, we show that the stabbing circle problem for parallel segments can be solved in $O(n \log^2 n)$ time and $O(n)$ space. A lower bound for the problem of computing a stabbing circle of minimum radius for parallel segments of equal length is shown in Section 6. Finally, in Section 7, we summarize and propose questions for future work.

2 Preliminaries and Definitions

In what follows, xx' denotes either a pair of points or a segment as convenient. For a pair x, y of points in the plane, let $d(x, y)$ denote the Euclidean distance between them, and let $bis(x, y)$ denote the perpendicular bisector of the segment xy. For a region $f \subset \mathbb{R}^2$, we denote its boundary as ∂f, and its closure as $\overline{f}$.

Definition 1. [16,23] *The* Hausdorff Voronoi diagram *of S is a partitioning of $\mathbb{R}^2$ into regions defined as follows:*

$$\mathsf{hreg}(aa') = \{p \in \mathbb{R}^2 \mid \forall bb' \in S \setminus \{aa'\}\colon \max\{d(p,a), d(p,a')\} < \max\{d(p,b), d(p,b')\}\};$$
$$\mathsf{hreg}(a) = \{p \in \mathsf{hreg}(aa') \mid d(p,a) > d(p,a')\}.$$

Note that $\mathsf{hreg}(a)$ and $\mathsf{hreg}(a')$ are subregions of $\mathsf{hreg}(aa')$ (see Figure 2a). Note as well, that $\mathsf{hreg}(aa')$ may have several connected components. Let $\mathsf{HVD}(S)$ denote the *graph structure* of the Hausdorff Voronoi diagram of S:

$$\mathsf{HVD}(S) = \mathbb{R}^2 \setminus \bigcup_{aa' \in S} (\mathsf{hreg}(a) \cup \mathsf{hreg}(a')).$$

An edge of $\mathsf{HVD}(S)$ is called *pure* if it separates the regions of two distinct segments, and it is called *internal* if it separates the subregions of the same segment. A vertex of $\mathsf{HVD}(S)$ is called *pure* if it is incident to three pure edges, and it is called *mixed* if it is incident to an internal edge. The pure vertices are defined by three distinct sites, and the mixed vertices by two distinct sites.[4]

Definition 2. [1,19] *The* farthest-color Voronoi diagram *is a partitioning of $\mathbb{R}^2$ into regions defined as follows:*

$$\mathsf{fcreg}(aa') = \{p \in \mathbb{R}^2 \mid \forall bb' \in S \setminus \{aa'\}\colon \min\{d(p,a), d(p,a')\} > \min\{d(p,b), d(p,b')\}\};$$
$$\mathsf{fcreg}(a) = \{p \in \mathsf{fcreg}(aa') \mid d(p,a) < d(p,a')\}.$$

The graph structure of this diagram is denoted as $\mathsf{FCVD}(S)$:

$$\mathsf{FCVD}(S) = \mathbb{R}^2 \setminus \bigcup_{aa' \in S} (\mathsf{fcreg}(a) \cup \mathsf{fcreg}(a')).$$

Similarly to the case of $\mathsf{HVD}(S)$, $\mathsf{fcreg}(a)$ and $\mathsf{fcreg}(a')$ are subregions of $\mathsf{fcreg}(aa')$, and $\mathsf{fcreg}(aa')$ may have several components. The edges and vertices of $\mathsf{FCVD}(S)$ are characterized as *pure* or *internal*, and *pure* or *mixed*, analogously to those of $\mathsf{HVD}(S)$ (see Figure 2b).

Let $\overline{\mathsf{hreg}}(\cdot)$ and $\overline{\mathsf{fcreg}}(\cdot)$ denote the closures of the respective regions.

When the segments in S are pairwise disjoint, the structural complexity of $\mathsf{HVD}(S)$ is $O(n)$ [16]. This particular case of pairwise disjoint segments has not been studied for $\mathsf{FCVD}(S)$. For arbitrary segments, the complexity of both diagrams is $O(n^2)$ [1,22].

[4] Any vertex of $\mathsf{HVD}(S)$ has degree three by the general position assumption that no four endpoints of segments in S are cocircular.

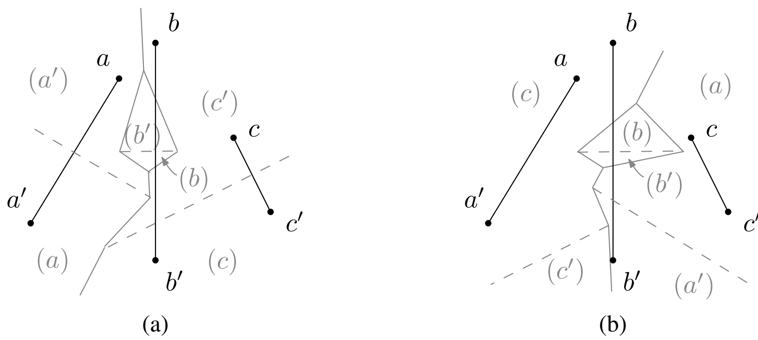

Fig. 2. (a) $\mathsf{HVD}(S)$, (b) $\mathsf{FCVD}(S)$. Pure and internal edges are represented in solid and dashed, respectively. The gray letters in parentheses label the respective regions.

Definition 3. *Given a point* p*, the* Hausdorff disk *of* p*, denoted* $D_h(p)$*, is the closed disk with center at* p *and radius* $d(p, a)$*, where* $p \in \overline{\mathsf{hreg}}(a)$*. The radius of* $D_h(p)$ *is called the* Hausdorff radius *of* p*, and is denoted as* $r_h(p)$*. The* farthest-color disk $D_f(p)$ *and the* farthest-color radius $r_f(p)$ *of* p *are defined analogously.*

The following lemma reveals the connection between the stabbing circle problem and the two cluster Voronoi diagrams, $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$.

Lemma 1. *Given a point* p*, there exists a stabbing circle centered at* p *if and only if* $r_f(p) < r_h(p)$*.*

Proof. Let c be a circle centered at p with radius r, and let D be the closed disk induced by c. Recall that c is a stabbing circle if and only if (1) each segment in S has an endpoint outside D; and (2) each segment in S has an endpoint in D. Condition (1) is equivalent to $r < r_h(p)$. Condition (2) is equivalent to $r_f(p) \leq r$. The claim follows. □

Now we are ready to define $\mathsf{FCVD}^\star(S)$, which is the closure of the locus of the centers of all stabbing circles for S.

Definition 4. *The* $\mathsf{FCVD}^\star(S)$ *is the locus of points in* $\mathbb{R}^2$ *for which the farthest-color radius is less than or equal to the Hausdorff radius, i.e.,* $\mathsf{FCVD}^\star(S) = \{p \in \mathbb{R}^2 : r_f(p) \leq r_h(p)\}$.

For any point on the boundary of $\mathsf{FCVD}^\star(S)$, its Hausdorff radius equals its farthest-color radius. Note that this equality also holds for some points in the interior of $\mathsf{FCVD}^\star(S)$. Any point p in the interior of $\mathsf{FCVD}^\star(S)$ such that $r_h(p) = r_f(p)$ lies on an internal edge of both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$, which separates centers of stabbing circles of two combinatorially different types, refer to Section 3.2 for more details.

By applying the transformation of Edelsbrunner and Seidel [18], both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ can be viewed as envelopes of *wedges* in 3D: Lift up the pairs of endpoints of the segments in S onto the unit paraboloid U, and join the lifted endpoints obtaining a set S' of segments in 3D. Each lifted endpoint a is mapped to a hyperplane tangent to U at point a. Thus, a segment in S is transformed to a pair of planes in 3D. The lower (resp., upper) envelope of such a pair forms a *lower* and (resp., *upper*) *wedge*, respectively. Then $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ correspond to the upper envelope of all lower wedges, and to the lower envelope of all upper wedges, respectively. This transformation is explicitly given for $\mathsf{HVD}(S)$ in Edelsbrunner et al. [16]. Thus, $\mathsf{FCVD}^{\star}(S)$ corresponds to the locus of points below $\mathsf{HVD}(S)$ and above $\mathsf{FCVD}(S)$.

The locus of points between two surfaces, defined by the above construction, is shown to be a representation of all combinatorially different stabbing planes for S' (if one exists) [16]. Further, the authors show that this locus is a set of $O(n^2)$ convex cells in 3D with $O(n^2)$ total complexity, and can be computed in $O(n^2)$ time and space.

We observe that a stabbing circle for S can be transformed into a stabbing plane for S' and vice versa.

We obtain the following result which, to the best of our knowledge, has not been explicitly stated anywhere before:

Theorem 1. *The stabbing circle problem for a set S of n arbitrary segments can be solved in $O(n^2)$ time and space.*

Claverol [10] showed that a set S of segments might have $\Theta(n^2)$ combinatorially different stabbing circles; see Figure 3. In the construction, each pair $\{a_i, a_j\}$ of points in the upper arc defines a stabbing circle that leaves the endpoints in the upper arc between a_i and a_j outside the circle. Hence, the $\Theta(n^2)$ stabbing circles defined in this way are combinatorially different. We remark that it is possible to perform a small perturbation so that the general position assumptions are satisfied.

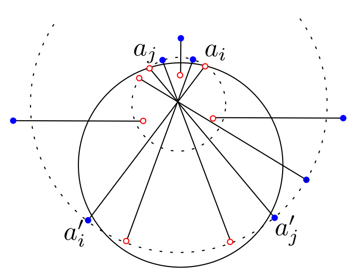

Fig. 3. A set with $\Theta(n^2)$ combinatorially different stabbing circles, and the stabbing circle defined by $\{a_i, a_j\}$.

3 Properties of $\mathsf{HVD}(S)$, $\mathsf{FCVD}(S)$, and $\mathsf{FCVD}^{\star}(S)$

In this section we investigate structural properties of the geometric structures involved in the stabbing circle problem. First, in Section 3.1, we list basic properties of the Hausdorff and the farthest-color Voronoi diagrams. These are later used to derive structural properties of $\mathsf{FCVD}^{\star}(S)$ and the correctness of our solution to the stabbing circle problem. The structure of the Hausdorff Voronoi diagram is well known, see e.g., [16,22,23];

however, this is not the case for the farthest-color diagram. We use Section 3.1 to derive some useful properties for $\mathsf{FCVD}(S)$. Then, in Section 3.2 we investigate the structure of $\mathsf{FCVD}^{\star}(S)$, and we characterize its faces and their complexity, linking them to features of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. We show that every face of $\mathsf{FCVD}^{\star}(S)$ corresponds to a unique combinatorially distinct solution of the stabbing circle problem. Finally, in Section 3.3, we complete the structural complexity analysis of $\mathsf{FCVD}^{\star}(S)$ and count its faces. The properties derived in this section are later used in Section 4 to obtain our algorithm that computes all the combinatorially distinct stabbing circles for S.

3.1 Properties of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$

We list some structural properties of the Hausdorff and the farthest-color Voronoi diagrams, which are used by our algorithms. First, a *visibility property* of both diagrams is summarized in the following lemma. For $\mathsf{HVD}(S)$, this property follows directly from [22, Property 2] (item (a)). We prove an equivalent property for $\mathsf{FCVD}(S)$ (item (b)). Items (a),(b) are illustrated in Figure 4a,b respectively.

Lemma 2. *Consider* $\mathsf{hreg}(a)$ *and* $\mathsf{fcreg}(a)$*, where* $aa' \in S$ *and* $|S| > 1$.

(a) For a point p *in* $\mathsf{hreg}(a)$*, the segment* ap *intersects* $\partial\mathsf{hreg}(a)$ *exactly once and the intersection point lies on an internal edge of* $\mathsf{HVD}(S)$*. For a point* q *on an internal edge of* $\partial\mathsf{hreg}(a)$*, the segment* aq *does not intersect* $\mathsf{hreg}(a)$.

(b) For a point p *in* $\mathsf{fcreg}(a)$*, the segment* ap *intersects* $\partial\mathsf{fcreg}(a)$ *exactly once and the intersection point lies on a pure edge of* $\mathsf{FCVD}(S)$*. For a point* q *on a pure edge of* $\partial\mathsf{fcreg}(a)$*, the segment* aq *does not intersect* $\mathsf{fcreg}(a)$.

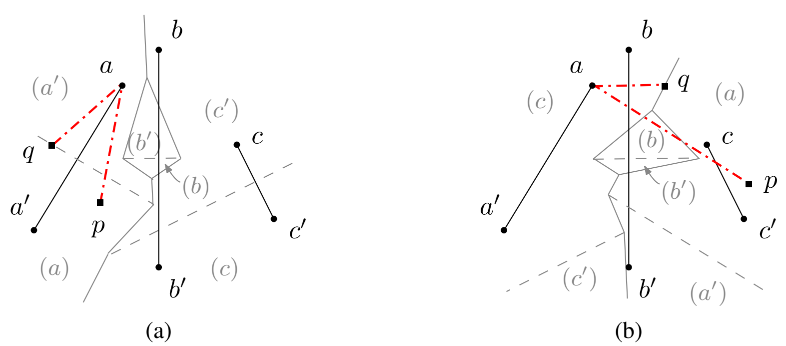

Fig. 4. Illustration for the statement of Lemma 2 based on the diagrams from Figure 2: line segments ap and aq are shown in red dash-dotted lines.

Proof. We only prove item (b). If $|S| > 1$, $a \notin \overline{\mathsf{fcreg}}(a)$, thus, ap intersects $\partial\mathsf{fcreg}(a)$ at least once. Suppose ap intersects $\partial\mathsf{fcreg}(a)$ more than once, and let x and y be the first two intersection points that are encountered when moving from p to a. Points x and

y are on $\partial\mathsf{fcreg}(a)$ and the segment xy is outside $\mathsf{fcreg}(a)$. Thus, $D_f(x)$ and $D_f(y)$ contain at least one endpoint of every segment in S, their boundary passes through a, and a' is outside both disks. Let w be a point on xy and let $D(w)$ be the closed disk centered at w whose boundary passes through a. Since the distance of any point on the line through x, y from a increases as we move away from a, $D_f(y) \subset D(w) \subset D_f(x)$. But then $D(w)$ must contain all points in $D_f(y)$, $a' \notin D(w)$, and $a \in \partial D(w)$. Thus, $D(w) = D_f(w)$; hence, $w \in \mathsf{fcreg}(a)$. We obtain a contradiction.

The second claim in item (b) follows directly from the first one by considering a point $p \in \mathsf{fcreg}(a)$ that is infinitesimally close to q. □

Figure 2 illustrates a Hausdorff and a farthest-color Voronoi diagram. Every component of a Hausdorff region $\mathsf{hreg}(aa')$ contains exactly one internal edge [22, Property 3] (see the dashed lines in Figure 2a). In the following lemmas we show properties of a similar nature for $\mathsf{FCVD}(S)$.

Lemma 3. *Any bounded face of* $\mathsf{fcreg}(aa')$ *contains an internal edge. (See the faces of region* $\mathsf{fcreg}(bb')$ *in Figure 2b.)*

Proof. Let f be a bounded face of $\mathsf{fcreg}(aa')$. Suppose for the sake of contradiction that f contains no internal edge. Then f is a face of $\mathsf{fcreg}(a)$, for an endpoint a of aa', such that ∂f consists solely of pure edges. Let p be a point in f and consider the ray r from a through p. Let q be the first intersection point between r and ∂f, as we move on r starting at p away from a (such point exists because f is bounded). Since f consists solely of pure edges, q is a point on a pure edge and $pq \in f$, yielding a contradiction to the last claim of Lemma 2b. □

Lemma 4. *For any face of* $\mathsf{fcreg}(a)$*, the portion of its boundary that is formed by pure edges is connected. (See the solid lines in Figure 2b.)*

Proof. Suppose $|S| > 1$ as otherwise the claim holds trivially. Let f be a face of $\mathsf{fcreg}(a)$. Clearly, any internal edge on ∂f is a portion of $bis(a, a')$. Since $|S| > 1$, $a \notin \mathsf{fcreg}(a)$, but a lies in the same halfplane induced by $bis(a, a')$ as $\mathsf{fcreg}(a)$. This implies that a lies in a region of the plane bounded by $bis(a, a')$ and one of the connected components of pure edges of ∂f (such regions are shown with tiling patterns in Figure 5). If ∂f contained more than one such connected components of pure edges, then any point x on any additional component would violate Lemma 2b. See Figure 5. □

The following property of $\mathsf{FCVD}(S)$ is only used in Section 5.

Lemma 5. $\mathsf{FCVD}(S)$ *has* $O(n)$ *unbounded faces.*

Proof. We observe that each unbounded face of $\mathsf{FCVD}(S)$ corresponds exactly to one face of the *farthest-segment Voronoi diagram* of S, $\mathsf{FsVD}(S)$. To make the correspondence one-to-one we remove the internal unbounded edges of $\mathsf{FCVD}(S)$ and merge the incident faces of $\mathsf{fcreg}(a)$ and $\mathsf{fcreg}(a')$, for every $aa' \in S$, into one face of $\mathsf{fcreg}(aa')$. Then a face f in $\mathsf{FCVD}(S)$ is unbounded in a direction ϕ if and only if a face f' in $\mathsf{FsVD}(S)$ is unbounded in the same direction ϕ. See [24] for the properties of the faces

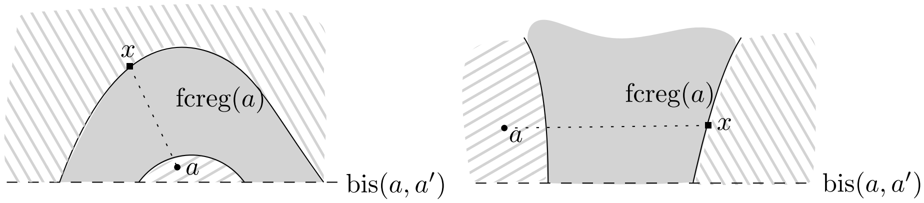

Fig. 5. Illustration for Lemma 4: two variants of an impossible situation.

of $\mathsf{FsVD}(S)$ at infinity, which are identical to those of $\mathsf{FCVD}(S)$. The total number of faces in $\mathsf{FsVD}(S)$ is $O(n)$ [4], thus, the same holds for the unbounded faces of $\mathsf{FCVD}(S)$. Note that a component of $\mathsf{fcreg}(aa')$ can contain at most two unbounded portions of $bis(a, a')$, thus, having removed the internal unbounded edges of $\mathsf{FCVD}(S)$ has no effect on the derived bound. □

3.2 Properties of $\mathsf{FCVD}^\star(S)$

In this section, we characterize the boundary of $\mathsf{FCVD}^\star(S)$, its connected components (for brevity, *components*), and its *faces*. We observe that the faces of $\mathsf{FCVD}^\star(S)$ are in one-to-one correspondence with the combinatorially different solutions for the stabbing circle problem for S. We characterize the faces of $\mathsf{FCVD}^\star(S)$ according to their intersection with $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$, and this characterization forms the basis of our algorithm to compute $\mathsf{FCVD}^\star(S)$ (and thus to solve the stabbing circle problem) presented in Section 4.

We proceed with describing the boundary of the components of $\mathsf{FCVD}^\star(S)$. Notice that a component is unbounded in a direction ϕ if and only if there exists a stabbing line for S that is orthogonal to ϕ.

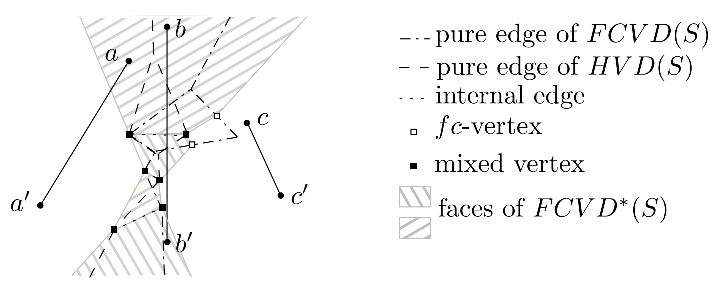

Fig. 6. $\mathsf{FCVD}^\star(S)$ where S is the set of segments $\{aa', bb', cc'\}$ from Figure 2.

We first observe that the vertices of $\partial\mathsf{FCVD}^\star(S)$ are all incident to edges of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$. Indeed, any vertex of $\partial\mathsf{FCVD}^\star(S)$ is caused by switching the region of either $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$ (in a traversal of $\partial\mathsf{FCVD}^\star(S)$), which happens exactly when $\partial\mathsf{FCVD}^\star(S)$ meets an edge of one of the diagrams. Thus we distinguish three types of vertices of $\partial\mathsf{FCVD}^\star(S)$: (1) vertices incident to pure edges of

$\mathsf{HVD}(S)$, called *h-vertices*; (2) vertices incident to pure edges of $\mathsf{FCVD}(S)$, called *fc-vertices*; and (3) vertices incident to internal edges of either diagram, called *mixed vertices*. See Figure 6 that illustrates faces of $\mathsf{FCVD}^{\star}(S)$ and vertices of $\partial\mathsf{FCVD}^{\star}(S)$. Notice that *h-vertices* and *fc-vertices* belong to the interior of corresponding pure edges. Moreover, the boundary of $\mathsf{FCVD}^{\star}(S)$ intersects these pure edges transversally, thus each h- or fc- vertex is adjacent to a portion of a pure edge of, respectively, $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ within $\mathsf{FCVD}^{\star}(S)$.

The following lemma implies that each mixed vertex of $\partial\mathsf{FCVD}^{\star}(S)$ is a mixed vertex of either $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, hence, the name for such vertices. Moreover, this lemma and its corollary lead to the definition of a *face* of $\mathsf{FCVD}^{\star}(S)$, which will then be treated as an atomic piece of $\mathsf{FCVD}^{\star}(S)$.

Lemma 6. *Point $p \in bis(a, a')$ is in $\mathsf{FCVD}^{\star}(S)$ if and only if p lies on an internal edge of both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. If p is on $\partial\mathsf{FCVD}^{\star}(S)$, then p is a mixed vertex of either $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$.*

Proof. Since $p \in bis(a, a')$, either both a, a' are outside $D_h(p)$ or they lie on its boundary. Symmetrically, either both a, a' are in the interior of $D_f(p)$ or on its boundary.

Suppose $p \in \mathsf{FCVD}^{\star}(S)$. Then $D_f(p) \subseteq D_h(p)$. By the above argument, $D_f(p) = D_h(p)$ and both a, a' lie on the boundary of this disk. Therefore, $r_h(p) = r_f(p) = pa = pa'$, and the claim follows.

Suppose that p lies on an internal edge of both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ that separates the respective regions of a and of a'. Since p lies on such edge of $\mathsf{HVD}(S)$, $r_h(p) = pa$, and since p lies on such edge of $\mathsf{FCVD}(S)$, $r_f(p) = pa$. Thus $r_h(p) = r_f(p)$, hence, $p \in \mathsf{FCVD}^{\star}(S)$.

Now we prove the second part of the statement. Since $p \in bis(a, a')$ is on $\partial\mathsf{FCVD}^{\star}(S)$, $D_h(p) = D_f(p)$, and the boundary of this disk passes through a, a'. It is easy to see that this boundary also passes through an endpoint c of some segment $cc' \in S$ (otherwise, the center of this disk could move in any direction while still being in $\mathsf{FCVD}^{\star}(S)$, contradicting the fact that $p \in \partial\mathsf{FCVD}^{\star}(S)$). If the other endpoint c' of cc' is inside this disk, then p is a mixed vertex of $\mathsf{HVD}(S)$, and if it is outside, p is a mixed vertex of $\mathsf{FCVD}(S)$. $\square$

Corollary 1. *No mixed vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ may lie in the interior of $\mathsf{FCVD}^{\star}(S)$.*

Proof. Suppose p is a mixed vertex of $\mathsf{HVD}(S)$ such that $p \in \mathsf{FCVD}^{\star}(S)$. Let e denote the internal edge of $\mathsf{HVD}(S)$ incident to p; e is a portion of $bis(a, a')$, for some $aa' \in S$. Since p is a vertex of e, there is a point p' on $bis(a, a')$ infinitesimally close to p such that p' does not belong to e. Observe that $r_h(p') < p'a$, and both a, a' are outside the Hausdorff disk of p'. This implies that p' is not in $\mathsf{FCVD}^{\star}(S)$, and therefore p is not in the interior of $\mathsf{FCVD}^{\star}(S)$. A similar argument (with $r_f(p') > p'a$) proves the case of p being a mixed vertex of $\mathsf{FCVD}(S)$. $\square$

A component of $\mathsf{FCVD}^{\star}(S)$ may (or may not) contain *internal edges* that partition the component into disjoint open *faces*. An internal edge of $\mathsf{FCVD}^{\star}(S)$ is the common portion of one internal edge of $\mathsf{HVD}(S)$ and one such edge of $\mathsf{FCVD}(S)$ within a component of $\mathsf{FCVD}^\star(S)$. A component of $\mathsf{FCVD}^\star(S)$ that contains internal edges is called a *multiple-face component* as it consists of more than one face. A component that contains no internal edges is called *single-face component* and its interior is a single face of $\mathsf{FCVD}^\star(S)$. To compute $\mathsf{FCVD}^\star(S)$ we need to identify its components. In the following we show that internal edges of $\mathsf{FCVD}^\star(S)$ are always incident to mixed vertices of either $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, and thus, the multiple-face components of $\mathsf{FCVD}^\star(S)$ can be easily identified. We distinguish the single-face components in two types: those that contain a vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$; and those that contain no such vertex, which are of a special form (see Lemma 11). Identifying the latter type of single-face components poses a major difficulty to our algorithm.

The following lemma shows that the faces of $\mathsf{FCVD}^\star(S)$ correspond exactly to combinatorially different stabbing circles. Thus, each single-face component reveals exactly one combinatorially distinct solution for the stabbing circle problem; the multiple-face components reveal a number of such solutions, exactly one for each of their faces.

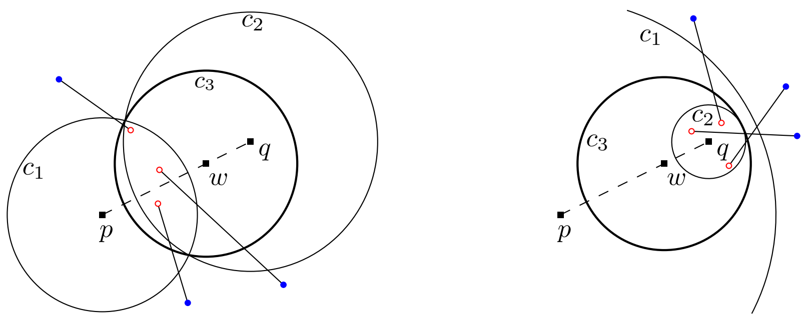

(a) (b)

Fig. 7. Illustration for the proof of Lemma 7.

Lemma 7. *Two stabbing circles are combinatorially different if and only if their centers lie in different faces of* $\mathsf{FCVD}^\star(S)$.

Proof. Let c_1 and c_2 be two combinatorially different stabbing circles, and let respectively p and q be their centers. There is a segment $aa' \in S$ such that a is enclosed in c_1, and a' is enclosed in c_2. Observe that p and q lie in different halfplanes with respect to $bis(a, a')$. If p and q were in the same face f of $\mathsf{FCVD}^\star(S)$, then there would be a path π connecting p and q such that π lies entirely in f. Since p and q are separated by $bis(a, a')$, path π would need to cross $bis(a, a')$ at a point t. However, point t cannot lie inside f due to Lemma 6, which would derive a contradiction.

Next, suppose that the stabbing circles c_1 and c_2 are combinatorially equivalent. Let w be a point on the segment pq. We show that w is the center of a stabbing circle c_3 that is combinatorially equivalent to c_1 and c_2. Indeed, if c_1 and c_2 intersect in two points, then let c_3 be the circle centered at w and passing through these two intersection points; see Figure 7a. If c_2 is enclosed in c_1, let c_3 be the minimum circle centered at w and enclosing c_2; see Figure 7b. Let D_1, D_2 and D_3 be the disks corresponding to c_1, c_2 and c_3 respectively. Observe that $D_1 \cap D_2 \subset D_3$ and $\big((\mathbb{R}^2 \setminus D_1) \cap (\mathbb{R}^2 \setminus D_2)\big) \subset (\mathbb{R}^2 \setminus D_3)$. Thus all the endpoints of segments in S that are enclosed in c_1 and c_2 are also enclosed in c_3, and the ones that lie outside of c_1 and c_2 also lie outside of c_3. This implies that c_3 is a stabbing circle that is combinatorially the same as c_1 and c_2. By Lemma 1, $r_f(w) < r_h(w)$, and thus w is in the interior of $\mathsf{FCVD}^\star(S)$ and does not lie on an internal edge of it. This implies that the closed segment pq lies in one face of $\mathsf{FCVD}^\star(S)$. $\square$

The second part of the above proof implies that, for any pair of points p, q within one face f of $\mathsf{FCVD}^\star(S)$, the segment pq lies in f. This proves the following property:

Corollary 2. *The faces of* $\mathsf{FCVD}^\star(S)$ *are convex.*

$\mathsf{FCVD}^\star(S)$ has the following visibility property, which is used in the proofs of subsequent lemmas in this section.

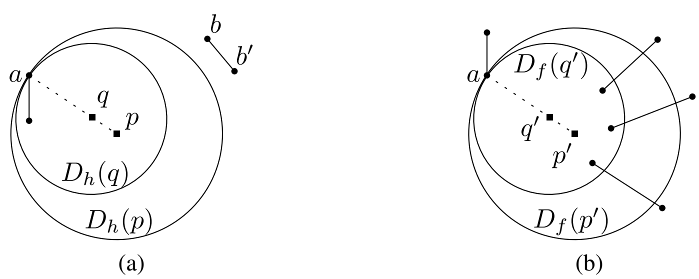

Fig. 8. Illustration for the proof of Lemma 8.

Lemma 8. *(a) Let p be a point outside $\mathsf{FCVD}^\star(S)$, and let aa' be a segment in S such that $p \in \overline{\mathsf{hreg}}(a)$. Then the entire segment $(pa \cap \overline{\mathsf{hreg}}(a))$ is outside $\mathsf{FCVD}^\star(S)$.*
(b) Let p' be a point in $\mathsf{FCVD}^\star(S)$, and let aa' be a segment in S such that $p' \in \overline{\mathsf{fcreg}}(a)$. Then the entire segment $(p'a \cap \overline{\mathsf{fcreg}}(a))$ is in $\mathsf{FCVD}^\star(S)$.

Proof. By Lemma 2, $(pa \cap \overline{\mathsf{hreg}}(a))$ and $(p'a \cap \overline{\mathsf{fcreg}}(a))$ are segments.

(a) Let q be a point on $pa \cap \overline{\mathsf{hreg}}(a)$; see Figure 8a. Clearly, $D_h(q) \subseteq D_h(p)$. Since $p \notin \mathsf{FCVD}^\star(S)$, it follows that $r_f(p) > r_h(p)$ and, in consequence, there exists a segment $bb' \in S$ such that $b, b' \notin D_h(p)$. Since $D_h(q) \subseteq D_h(p)$, points b and b' are also outside $D_h(q)$. Thus, $q \notin \mathsf{FCVD}^\star(S)$.

(b) Let q' be a point in $p'a \cap \overline{\mathsf{fcreg}}(a)$; see Figure 8b. Since $p' \in \mathsf{FCVD}^\star(S)$, we have that $r_f(p') \leq r_h(p')$ and thus each segment from S has at least one endpoint outside $D_f(p')$ or on its boundary. Since $D_f(q') \subset D_f(p')$, the same holds for $D_f(q')$. Thus, $D_f(q') \subset D_h(q')$, that is, $q' \in \mathsf{FCVD}^\star(S)$. $\square$

The following two lemmas analyze the intersection of a face of $\mathsf{FCVD}^\star(S)$ with $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. In particular, we characterize connectedness and non-emptiness of such intersections. These are properties that enable us to efficiently identify the faces of $\mathsf{FCVD}^\star(S)$.

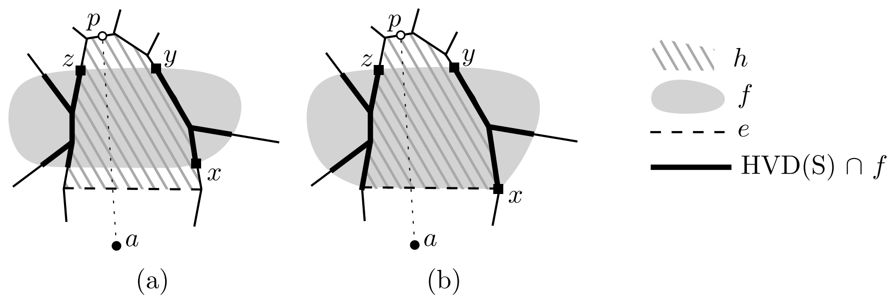

Fig. 9. A face f of $\mathsf{FCVD}^\star(S)$ in an impossible situation where $f \cap \mathsf{HVD}(S)$ is disconnected. Two variants with respect to the internal edge e: (a) e is outside $\overline{f}$; (b) e is on ∂f.

Lemma 9. *For a face f of $\mathsf{FCVD}^\star(S)$, both $f \cap \mathsf{HVD}(S)$ and $f \cap \mathsf{FCVD}(S)$ are connected.*

Proof. Suppose for the sake of contradiction that $f \cap \mathsf{HVD}(S)$ is disconnected, see Figure 9. Then there is a face h of $\mathsf{hreg}(a)$ such that $f \cap \partial h$ has at least two connected components. By [22, Property 3], ∂h contains exactly one internal edge e, $e \subseteq bis(a, a')$. By the definition of a face of $\mathsf{FCVD}^\star(S)$, e cannot intersect the interior of f; it can only border f or it must be outside $\overline{f}$, see Figure 9a and b.

Consider the first connected component of $f \cap \partial h$ that follows e in a counterclockwise traversal of ∂h. Let the endpoints of this component be called x and y (in the order of the traversal), see Figure 9. Let z be the first point of f encountered as we continue traversing ∂h counterclockwise beyond y. That is, z is a point in another connected component of $f \cap \partial h$. Consider the portions of ∂f and ∂h respectively from point y to point z. The portion of ∂f from y to z is inside h. The portion of ∂h from y to z consists solely of pure edges. Thus a point p in this portion of ∂h infinitesimally close to y is outside $\mathsf{FCVD}^\star(S)$. Then by Lemma 8a $pa \cap h$ is outside of $\mathsf{FCVD}^\star(S)$. By Lemma 2a, the segment pa intersects $\partial\mathsf{hreg}(a)$ once and the intersection point lies on e. Thus, $(pa \cap h) \cap f$ is not empty. We obtain a contradiction.

The proof for $f \cap \mathsf{FCVD}(S)$ is similar. In particular, if $f \cap \mathsf{FCVD}(S)$ is disconnected, then $f \cap \partial h$ is disconnected, for a face h of $\mathsf{fcreg}(a)$. Similarly to the above, by Lemma 4 and Lemma 2b, there is a point p on a pure edge on $\partial\mathsf{fcreg}(a)$ such that $p \notin \mathsf{FCVD}^\star(S)$, and the supporting line of the segment pa intersects f inside $\mathsf{fcreg}(a)$. Let p' be a point in this intersection. Then point $p \in p'a$, $p \in \overline{\mathsf{fcreg}}(a)$ and $p \notin \mathsf{FCVD}^\star(S)$, which contradicts Lemma 8b. □

Lemma 10. *Let f be a bounded face of $\mathsf{FCVD}^\star(S)$. (a) At least one of $f \cap \mathsf{HVD}(S)$ and $f \cap \mathsf{FCVD}(S)$ must be non-empty. (b) If one of $f \cap \mathsf{HVD}(S)$ or $f \cap \mathsf{FCVD}(S)$ is empty, then f belongs to a component of $\mathsf{FCVD}^\star(S)$ with multiple faces.*

Proof. To prove item (a), recall that any vertex on the boundary of a component of $\mathsf{FCVD}^\star(S)$ is incident to an edge of $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$ within that component.

Fig. 10. Illustration for the proof of Lemma 10. An impossible situation where a face f of $\mathsf{FCVD}^{\star}(S)$ (shaded) is such that (a) $\mathsf{HVD}(S) \cap f$ is empty; (b) $\mathsf{FCVD}(S) \cap f$ is empty.

Recall also that a face of $\mathsf{FCVD}^{\star}(S)$ is an open region that does not include internal edges of the diagrams, but can be incident to portions of them. Therefore any h- or fc-vertex on the boundary of a face of $\mathsf{FCVD}^{\star}(S)$ is incident to a pure edge of respectively $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ within this face. For the mixed vertices on the boundary of the face it might not hold, since the internal edges incident to them lie on the bounday of the face of $\mathsf{FCVD}^{\star}(S)$ by definition.

Suppose for the sake of contradiction that $f \cap \mathsf{HVD}(S) = \emptyset$ and $f \cap \mathsf{FCVD}(S) = \emptyset$. Then all the vertices of $\mathsf{FCVD}^{\star}(S)$ on the boundary of f are mixed vertices. Since f is bounded and non-empty, the number of vertices on ∂f is at least four. All these mixed vertices cannot lie on the bisector of the same segment in S because f is convex (see Corollary 2). Therefore f is incident to a segment of the internal edges of at least two regions of $\mathsf{HVD}(S)$. Let points s, t be respectively a point in the first segment and a point in the second one. Points s and t belong to the Hausdorff Voronoi regions of two different segments, thus the open segment st crosses at least one pure edge of $\mathsf{HVD}(S)$. Since f is convex, the entire open segment st is contained in f, and thus $f \cap \mathsf{HVD}(S)$ is not empty. We arrive to a contradiction.

Now we prove item (b). Suppose for the sake of contradiction that f is bounded, $f \cap \mathsf{HVD}(S)$ is empty, and $\overline{f}$ is a single-face component of $\mathsf{FCVD}^{\star}(S)$. Since $f \cap \mathsf{HVD}(S)$ is empty, f is entirely contained in $\mathsf{hreg}(a)$ for an endpoint a of some segment $aa' \in S$. Since segments in S do not share endpoints, no pure edge of $\mathsf{HVD}(S)$ can overlap with an edge of $\partial\mathsf{FCVD}^{\star}(S)$ (they can only intersect in one point). Then there is a point $p \in \mathsf{hreg}(a) \setminus f$ such that pa intersects f and $p \notin \mathsf{FCVD}^{\star}(S)$ (see Figure 10a). But by Lemma 8a, since $p \notin \mathsf{FCVD}^{\star}(S)$ the entire segment $(pa \cap \overline{\mathsf{hreg}}(a))$ must be outside $\mathsf{FCVD}^{\star}(S)$. We obtain a contradiction.

The symmetric statement about $f \cap \mathsf{FCVD}(S)$ can be shown as follows. Suppose that $f \subset \mathsf{fcreg}(a)$, see Figure 10b. Similarly to the above, edges of ∂f and $\partial\mathsf{fcreg}(a)$ do not overlap. Pick a point $p' \in f$. By Lemma 8b, the entire segment $p'a \cap \mathsf{fcreg}(a)$ lies in f. Recall that $a \notin \overline{\mathsf{fcreg}}(a)$ (since segments in S do not share endpoints), and thus f must intersect $\mathsf{FCVD}(S)$. We obtain a contradiction. $\square$

Finally, in the next lemma, we explore a special type of single-face components of $\mathsf{FCVD}^{\star}(S)$: the ones that contain no vertices of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$. The faces that correspond to bounded components of this type create the main difficulty in computing $\mathsf{FCVD}^{\star}(S)$.

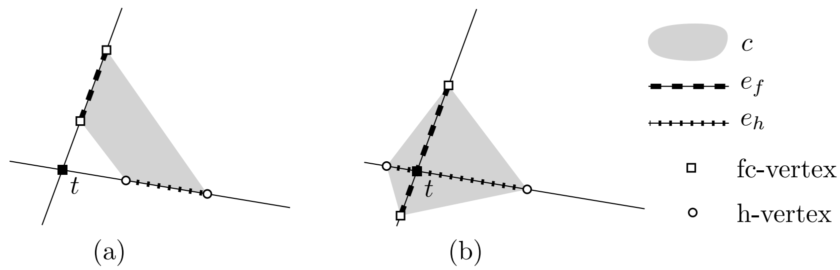

Fig. 11. Illustration for the proof of Lemma 11. A bounded component c of $\mathsf{FCVD}^{\star}(S)$ that does not contain a vertex of $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$: (a) an impossible situation; (b) the only possible situation.

Lemma 11. *Let c be a component of $\mathsf{FCVD}^{\star}(S)$ that contains no vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$. If c is bounded, then c contains exactly one intersection of two pure edges (one of $\mathsf{HVD}(S)$ and one of $\mathsf{FCVD}(S)$), and ∂c is a quadrilateral. If c is unbounded, then c contains an unbounded portion of a pure edge of either $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$.*

Proof. Since c does not contain any mixed vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, by Lemma 6, c does not intersect any internal edge of the diagrams. Thus ∂c may only contain h- and fc-vertices (no mixed vertices). Further, the interior of c is a single face f of $\mathsf{FCVD}^{\star}(S)$ ($\overline{f} = c$).

Suppose first that c is bounded. Lemmas 9 and 10 imply that $f \cap \mathsf{HVD}(S)$ is, respectively, connected and non-empty. Together with the fact that f contains no vertices of $\mathsf{HVD}(S)$, this implies that $f \cap \mathsf{HVD}(S)$ is an (open) line segment, whose endpoints are two h-vertices on ∂c. By the analogous argument for $f \cap \mathsf{FCVD}(S)$, ∂c has exactly two fc-vertices. Therefore $\partial c = \partial f$ is a quadrilateral. See Figure 11b.

Let e_h and e_f denote respectively the open line segments $f \cap \mathsf{HVD}(S)$ and $f \cap \mathsf{FCVD}(S)$; see Figure 11. Let t be the point of intersection between the supporting lines of e_h and e_f (the case where these lines are parallel is similar). We will show that $t \in f$. Note that it is impossible that an h-vertex is followed by the other h-vertex on ∂c (see Figure 11a), as it would imply that e_h and e_f lie entirely on ∂c, thus $f \cap \mathsf{HVD}(S)$ would be empty (as well as $f \cap \mathsf{FCVD}(S)$), what contradicts Lemma 10. Thus the different type of vertices interleave on ∂f (see Figure 11b), which implies that $t \in f$. This completes the proof of the first statement.

Now suppose that c is unbounded. Note that ∂c consists of at least two edges, and c is convex (see Corollary 2). Thus there are two unbounded edges on ∂c that have different supporting lines. Denote these edges as e and g. If c contained no unbounded portions of pure edges of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, then both e and g would be contained in $\mathsf{hreg}(a)$ and in $\mathsf{fcreg}(b)$, for some endpoints a and b of segments in S. But then both e and g must be portions of $bis(a, b)$. We obtain a contradiction. $\square$

We have explored different types of components of $\mathsf{FCVD}^{\star}(S)$, and we saw that a component may or may not be comprised of multiple faces, may or may not contain a vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, and it may or may not be bounded. The cumulative complexity of all components that contain multiple faces, or contain a vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, or are unbounded, is $O(|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$ (see the proof of Theorem 2 in Section 3.3). However, the bounded single-face components that do not contain a vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ (i.e., the quadrilateral components of Lemma 11) do not fall under this bound. We bound their number in the next section. Identifying these components poses the main challenge to our algorithm in Section 4.

3.3 Complexity of $\mathsf{FCVD}^{\star}(S)$

With some abuse of notation, we denote by $|\mathsf{HVD}(S)|$, $|\mathsf{FCVD}(S)|$ and $|\mathsf{FCVD}^{\star}(S)|$ respectively the number of edges of $\mathsf{HVD}(S)$, $\mathsf{FCVD}(S)$, and $\partial\mathsf{FCVD}^{\star}(S)$. We aim to connect $|\mathsf{FCVD}^{\star}(S)|$ with $|\mathsf{HVD}(S)|$ and $|\mathsf{FCVD}(S)|$.

We classify the segments in S with respect to a portion of a fixed pure edge of $\mathsf{HVD}(S)$. Let e be a connected portion of pure edge of $\mathsf{HVD}(S)$ that separates $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$, for two segments $aa', bb' \in S$. In particular, $e \subseteq bis(a,b)$. For the rest of this section, it is convenient to perform a rotation of the coordinate system so that e is horizontal. Let u and respectively v be the left and right endpoints of e.

If u is a mixed vertex of $\mathsf{HVD}(S)$, we redefine u as a point on e infinitesimally to the right, so that u is in the boundary of only $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$. We proceed analogously with v.

The Hausdorff disks $D_h(u)$ and $D_h(v)$ have a, b on the boundary, they contain aa', bb', and they do not contain any other segment of S. The same holds for the Hausdorff disk of any point of e. Hence, every segment $cc' \in S \setminus \{aa', bb'\}$ can be classified as follows (see Figure 12, left):

- cc' is of type *out* if both c and c' are outside $D_h(u) \cup D_h(v)$;
- cc' is of type *in* if either c or c' is contained in $D_h(u) \cap D_h(v)$ and the other endpoint is outside $D_h(u) \cup D_h(v)$;
- cc' is of type *left* if either c or c' is contained in $D_h(u) \setminus D_h(v)$ and the other endpoint is outside $D_h(u) \cup D_h(v)$;
- cc' is of type *right* if either c or c' is contained in $D_h(v) \setminus D_h(u)$ and the other endpoint is outside $D_h(u) \cup D_h(v)$;
- cc' is of type *middle* if either c or c' is contained in $D_h(u) \setminus D_h(v)$ and the other endpoint is contained in $D_h(v) \setminus D_h(u)$.

For any point $p \in e$, we use p_ℓ and p_r to denote two points in e infinitesimally close to p and lying to the left and right of p, respectively. Additionally, let M^e_{bis} denote the set of segments $cc' \in S$ of type *middle* for edge e such that e intersects an internal edge of $\mathsf{FCVD}(S)$ separating $\mathsf{fcreg}(c)$ from $\mathsf{fcreg}(c')$. Notice that this internal edge is a portion of $bis(c, c')$. We also set $m^e_{bis} = |M^e_{bis}|$.

Let m_{bis} denote the total number of pairs formed by a pure edge e of $\mathsf{HVD}(S)$ and a segment $cc' \in S$ such that $cc' \in M^e_{bis}$.

Lemma 12. *Let e be a portion of a pure edge of $\mathsf{HVD}(S)$. The number of faces of $\mathsf{FCVD}^{\star}(S)$ intersected by e is at most $1 + m^e_{bis}$.*

Proof. Let e' be the left-most portion (if any) of the interior of e contained in $\mathsf{FCVD}^{\star}(S)$, and let w be its right endpoint. Since $w \in \mathsf{FCVD}^{\star}(S)$, each segment

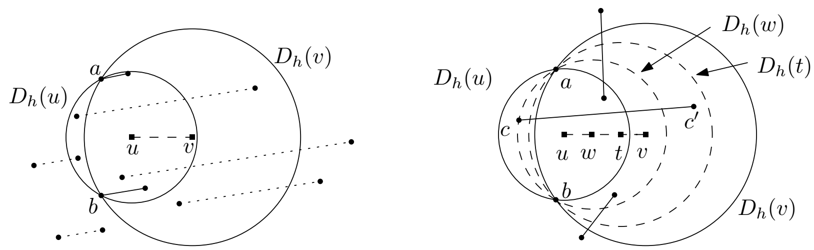

Fig. 12. Left: From top to bottom, the types of the dotted segments are *middle*, *left*, *in*, *right*, and *out*. Right: Illustration for the proof of Lemma 12.

from S has at least one endpoint inside $D_h(w)$. Let cc' be a segment in S such that $w_r \in \overline{\mathsf{fcreg}}(c)$. Since $w_r \notin \mathsf{FCVD}^\star(S)$, both c and c' are outside $D_h(w_r)$. This implies that one of c or c' (say, c) lies in the portion of the boundary of $D_h(w)$ contained in $D_h(u)$, and the other endpoint lies outside $D_h(w)$ (see Figure 12, right). Consequently, for any point t in $w_r v$, $D_h(t)$ does not contain c. Suppose that there are more portions of e contained in $\mathsf{FCVD}^\star(S)$, and let e'' be the left-most one. Then, for any point t in e'', $D_h(t)$ contains c' but not c. In particular, c' lies in $D_h(v) \setminus D_h(u)$ and cc' is of type *middle* for e. Hence, $bis(c, c')$ intersects e at a point q leaving e' to its left and e'' to its right. If q lies on an internal edge of $\mathsf{FCVD}(S)$ separating $\mathsf{fcreg}(c)$ from $\mathsf{fcreg}(c')$, then $cc' \in M^e_{bis}$ and we assign e'' to cc'.

Otherwise, q lies in $\overline{\mathsf{fcreg}}(d)$, for some $dd' \in S \setminus \{aa', bb', cc'\}$. Since $w_r \in \overline{\mathsf{fcreg}}(c)$, the segment $w_r q$ crosses at least one edge of $\mathsf{FCVD}(S)$. We start traversing the segment $w_r v$ starting from w_r. Suppose that we leave $\overline{\mathsf{fcreg}}(c)$ and we enter $\overline{\mathsf{fcreg}}(f)$, for some $ff' \in S$. Notice that this happens before we reach e''. Since $w \in \mathsf{FCVD}^\star(S)$, $w_r \notin \mathsf{FCVD}^\star(S)$ and $f \neq c'$, the point f does not lie in $D_h(u) \cap D_h(v)$, $D_h(v) \setminus D_h(u)$ or outside $D_h(u) \cup D_h(v)$. Hence, it lies in $D_h(u) \setminus D_h(v)$. Furthermore, since $e'' \subseteq \mathsf{FCVD}^\star(S)$, ff' is of type *middle* and, for any point t in e'', $D_h(t)$ contains f'. Hence, $bis(f, f')$ intersects e at a point q' leaving e' to its left and e'' to its right. If q' lies on an internal edge of $\mathsf{FCVD}(S)$ separating $\mathsf{fcreg}(f)$ from $\mathsf{fcreg}(f')$, then $ff' \in M^e_{bis}$ and we assign e'' to ff'. Otherwise, we continue traversing $w_r v$. The left endpoint of e'' is in the farthest-color region of a point in $D_h(v) \setminus D_h(u)$. Thus, as we traverse $w_r v$ and simultaneously $\mathsf{FCVD}(S)$, at some point we cross an edge of $\mathsf{FCVD}(S)$ separating the farthest-color region of a point in $D_h(u) \setminus D_h(v)$ from the farthest-color region of a point in $D_h(v) \setminus D_h(u)$. This is only possible when this edge is an internal edge of $\mathsf{FCVD}(S)$ separating the regions of two endpoints of a segment of type *middle*. We assign ee' to this segment in M^e_{bis}.

Next, we select the right endpoint of e'' and perform the same analysis. By repeating the same argument until we reach v, we obtain an assignment of the portions of the interior of e contained in $\mathsf{FCVD}^\star(S)$ (except for e') to segments in M^e_{bis}. Furthermore, in this assignment no segment is charged more than one portion. This completes the proof of the lemma. $\square$

Theorem 2. *Let S be a set of n segments in the plane in general position. Then $|\mathsf{FCVD}^\star(S)| = O(|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + m_{bis})$.*

Proof. Consider a face f of $\mathsf{FCVD}^\star(S)$. By Lemma 9, $f \cap \mathsf{HVD}(S)$ is a connected graph. Each vertex of this graph (if any) has degree three, since it is a vertex of $\mathsf{HVD}(S)$. By Euler's formula the number of h-vertices on ∂f is $O(H+1)$, where H is the number of vertices of $\mathsf{HVD}(S)$ inside f. Analogously, the number of fc-vertices on ∂f is $O(F+1)$, where F is the number of vertices of $\mathsf{FCVD}(S)$ inside f.

Any vertex on ∂f that is neither an h-vertex nor an fc-vertex is a mixed vertex of $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$. Clearly, each pure vertex of $\mathsf{HVD}(S)$ and of $\mathsf{FCVD}(S)$ lies in at most one face of $\mathsf{FCVD}^\star(S)$, and each mixed vertex of one of these diagrams lies on the boundary of exactly two faces of $\mathsf{FCVD}^\star(S)$. Thus the total complexity of the boundary of all components of $\mathsf{FCVD}^\star(S)$ that contain at least one vertex of $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$ is $O(\mathsf{HVD}(S) + \mathsf{FCVD}(S))$.

What remains is to bound the number of components of $\mathsf{FCVD}^\star(S)$ that do not contain a vertex of $\mathsf{HVD}(S)$ or of $\mathsf{FCVD}(S)$. We consider separately the unbounded and the bounded components of this type. The total number of unbounded ones is $O(|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$: By Lemma 11, each such unbounded component contains an unbounded portion of an edge of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, and clearly one edge can correspond to at most one component. By Lemma 11, each bounded component of $\mathsf{FCVD}^\star(S)$ with no vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ intersects exactly one pure edge of $\mathsf{HVD}(S)$, and this intersection is connected. Thus the number of such components of $\mathsf{FCVD}^\star(S)$ can be upper-bounded by the total number of intersections of all pure edges of $\mathsf{HVD}(S)$ with $\mathsf{FCVD}^\star(S)$. By Lemma 12, this is $O(|\mathsf{HVD}(S)| + m_{bis})$.

Summing up the total complexity of the boundary for all types of components of $\mathsf{FCVD}^\star(S)$ implies the claim. $\square$

4 Computing $\mathsf{FCVD}^\star(S)$

Recall from Section 3.2 that any unbounded component of $\mathsf{FCVD}^\star(S)$ corresponds to a stabbing line for S. All the stabbing lines and the corresponding components of $\mathsf{FCVD}^\star(S)$ can be found in $O(n \log n)$ time [15]. From now on we assume that S has no stabbing line, and therefore all components and faces of $\mathsf{FCVD}^\star(S)$ are bounded.

4.1 General algorithm

By Lemmas 6, 10 and 11, any bounded face f of $\mathsf{FCVD}^\star(S)$ has at least one of the following properties: (1) f is incident to a mixed vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, (2) f contains a pure vertex of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$, or (3) f contains exactly one segment of a pure edge of both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. Our algorithm, described in Figure 13, has three parts. The first part (Steps 1–7) computes the faces of property (1). The second part (Steps 8–12) computes the faces of property (2) that have not been computed so far. After that the faces that satisfy only property (3) are left, and they are computed in Steps 13–14.

We next make some remarks about the first two parts of the algorithm in Figure 13.

Algorithm COMPUTING $\mathsf{FCVD}^*(S)$

1. compute $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$;
2. **for** each non-marked mixed vertex v of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ **do**
3.   **if** $v \in \partial\mathsf{FCVD}^\star(S)$ **then**
4.   **for** each face f of $\mathsf{FCVD}^\star(S)$ incident to v;
5.   **if** f has not been visited yet **then**
6.   trace ∂f starting from v;
7.   mark all vertices of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ in $\overline{f}$;
8. **for** each non-marked pure vertex u of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ **do**
9.   **if** $u \in \mathsf{FCVD}^\star(S)$ **then**
10.   let f be the face of $\mathsf{FCVD}^\star(S)$ that contains u;
11.   find a point on ∂f and trace ∂f;
12.   mark all vertices of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ in $\overline{f}$;
13. **for** each segment e of a pure edge of $\mathsf{HVD}(S)$ that is outside the computed faces of $\mathsf{FCVD}^\star(S)$ **do**
14.   compute all faces of $\mathsf{FCVD}^\star(S)$ that intersect e;

Fig. 13. Algorithm to compute $\mathsf{FCVD}^\star(S)$.

Steps 3 and 9 are simple, after $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ are computed and preprocessed for point location queries. For example, if v is a vertex of say $\mathsf{HVD}(S)$, we first locate v in $\mathsf{FCVD}(S)$. Once we obtain a segment $aa' \in S$ such that $v \in \overline{\mathsf{fcreg}}(aa')$, we can compare the Hausdorff and farthest-color radii of v in constant time.

The loop in Step 4 performs exactly two iterations, since each mixed vertex on $\partial\mathsf{FCVD}^\star(S)$ has exactly two incident faces of $\mathsf{FCVD}^\star(S)$ (see Section 3.2).

In Step 11, given a point u inside face f, we proceed as follows to find a point on ∂f. We first locate u in both $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. Then we trace in both diagrams the vertical ray originating at u, until we reach the boundary of f. To perform this efficiently, we preprocess $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ for ray-shooting queries.

The third part of the algorithm (Steps 13–14) is discussed in Section 4.2, its correctness in Section 4.3, and its time complexity in Section 4.4 (specifically, see Theorem 3).

4.2 Searching in a pure edge of $\mathsf{HVD}(S)$

In this section, we present an algorithm to compute all faces of $\mathsf{FCVD}^\star(S)$ intersected by a given portion of a pure edge of $\mathsf{HVD}(S)$.

Due to the assumption that S does not have any stabbing line, all faces of $\mathsf{FCVD}^\star(S)$ we are searching for are bounded. Thus, the input portion of a pure edge of $\mathsf{HVD}(S)$ is assumed to be a line segment. The algorithm is given as a pseudocode in Figure 15.

Without loss of generality, we assume that the input segment $e = uv$ is horizontal, and that u is to the left of v. The algorithm assumes that, if e is shrunk infinitesimally from both sides (see Step 1), the resulting segment $u_r v_\ell$ has both endpoints outside $\mathsf{FCVD}^\star(S)$. To proceed with the description of the algorithm, we need to introduce some notation. For convenience, we assume that the segment uv is already shrunk.

Let $e = uv$ be a line segment on a pure edge of $\mathsf{HVD}(S)$ separating $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$, such that u and v are outside $\mathsf{FCVD}^{\star}(S)$. Additionally, if the segment ab intersects the interior of e, this intersection divides e into two portions, which we process separately. Note that neither u nor v are mixed vertices of $\mathsf{HVD}(S)$ (since uv is the result of shrinking a portion of a pure edge of $\mathsf{HVD}(S)$ from both sides).

We classify the segments in $S \setminus \{aa', bb'\}$ with respect to uv as segments of types *left*, *right*, *middle*, *in*, *out*, as in Section 3.3 (see Figure 12). Using this classification, for any point w in e, we define $type(w)$ as a set containing one element per each $cc' \in S$ such that $w \in \overline{\mathsf{fcreg}}(cc')$. The elements of $type(w)$ are defined as follows: Let cc' be a segment in S such that $w \in \overline{\mathsf{fcreg}}(cc')$. If cc' is not of type *middle*, then we add the type of cc' to $type(w)$. If cc' is of type *middle*, then either c or c' (say, c) is contained in $D_h(u) \setminus D_h(v)$, and the other endpoint (c') is contained in $D_h(v) \setminus D_h(u)$. We further differentiate the classification *middle* as follows: If w lies on $bis(c, c')$, then $mm \in type(w)$. Otherwise, if $w \in \overline{\mathsf{fcreg}}(c)$, then $ml \in type(w)$; if $w \in \overline{\mathsf{fcreg}}(c')$, then $mr \in type(w)$. When we need to specify cc', we do as follows: Imagine that $w \in \overline{\mathsf{fcreg}}(cc')$ and cc' is of type in. Then we say $in \in type(w)$ *caused by* cc'.

Further, we use $\tilde{l}$ to denote types *left* and ml, and we use $\tilde{r}$ to denote *right* and mr.

Definition 5. *A point w in e is a* changing point *if $\{\tilde{r}, \tilde{l}\} \subseteq type(w)$ (see Figure 14).*

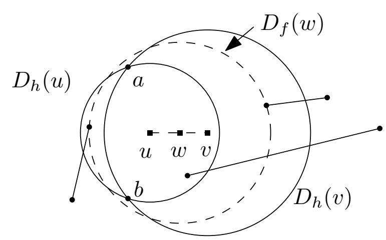

Fig. 14. w is a changing point (not in $\mathsf{FCVD}^{\star}(S)$).

A changing point might or might not be in $\mathsf{FCVD}^{\star}(S)$. Note that a changing point w is an intersection point between a pure edge of $\mathsf{HVD}(S)$ and a pure edge of $\mathsf{FCVD}(S)$. Intuitively, at w the point giving the farthest-color radius changes from being in $D_h(v) \setminus D_h(u)$ to being in $D_h(u) \setminus D_h(v)$, i.e., $\tilde{r} \in type(w_\ell)$ and $\tilde{l} \in type(w_r)$ (see Figure 14). Then we have the following.

Lemma 13. *(a) If e intersects a face of $\mathsf{FCVD}^{\star}(S)$ in a segment e', then there exists a point w in e' such that $in \in type(w)$ or w is a changing point. (b) If there is a point $w \in e$ such that $out \in type(w)$, then e does not intersect $\mathsf{FCVD}^{\star}(S)$. (c) If there is a point $w \in e$ such that $in \in type(w)$, then $w \in \mathsf{FCVD}^{\star}(S)$.*

Proof. Even though item (a) is the base of the algorithm SEARCH IN e, it is actually not needed to prove its correctness, given in Lemma 16, because Lemma 16, together with Lemmas 14 and 15, in a way reproves item (a). Therefore, we omit the proof of this item.

(b) Let $e = uv$, and let cc' be a segment in S that causes *out* in $type(w)$. Then both c and c' are outside $D_h(u) \cup D_h(v)$. For any point $x \in e$, $D_h(x) \subset D_h(u) \cup D_h(v)$, thus both c and c' are outside $D_h(x)$, and therefore $x \notin \mathsf{FCVD}^\star(S)$. This implies that $e \cap \mathsf{FCVD}^\star(S) = \emptyset$.

(c) If $in \in type(w)$, then there is a segment $cc' \in S$ such that $w \in \overline{\mathsf{fcreg}}(cc')$ and one of the endpoints of cc' is in $D_h(u) \cap D_h(v)$. This endpoint is the closest one to w among c, c'. Thus it is on the boundary of $D_f(w)$, and $D_f(w) \subset D_h(w)$, which implies the claim. □

Thus, it is enough to examine the changing points of e, and the points w such that $in \in type(w)$. To find such points, we use the *find-change query* subroutine, defined next. When dealing with a subsegment ts of e, or a pair (t, s) of points in e, we write the left-most point first.

Definition 6 (Find-change query). *The input of the query is a pair (t, s) of points in e, such that $type(t)$ contains $\tilde{r}$ but not $\tilde{l}$, and $type(s)$ contains $\tilde{l}$ but not $\tilde{r}$. The query returns a point w in the segment ts such that one of the following holds: (i) w is a changing point; (ii) $in \in type(w)$; (iii) $out \in type(w)$.*

For the sake of clarity, we defer the proof that the find-change query is well-defined, as well as the proof of correctness of the remaining pieces of the algorithm, to the next subsection.

Algorithm SEARCH IN $e = uv$
$uv \leftarrow u_r v_\ell$; (* shrink the segment uv infinitesimally *)
if $\tilde{r}$ is the only element in $type(u)$ **and** $\tilde{l}$ is the only element in $type(v)$ **then**
    perform a find-change query on (u, v);
    let w be the point returned by the query;
    if $out \in type(w)$ **then**
        return;
    if $w \in \mathsf{FCVD}^\star(S)$ **then**
        trace ∂f, where f is the face of $\mathsf{FCVD}^\star(S)$ that contains w;
        let q, q' be the points of $\partial f \cap uv$, where q is to the left of q';
    else (* w is a changing point in uv **and** $w \notin \mathsf{FCVD}^\star(S)$ *)
        let q, q' both be w;
    SEARCH IN uq;
    SEARCH IN $q'v$;
    return;
else return;

Fig. 15. Algorithm to compute all faces of $\mathsf{FCVD}^\star(S)$ intersected by $e = uv$.

We are now ready to describe our algorithm SEARCH IN e, which computes all the faces of $\mathsf{FCVD}^\star(S)$ intersected by e. The algorithm is illustrated in Figure 15; it uses the characterization from Lemma 13. At any time, the algorithm processes a subsegment of e. It first shrinks e infinitesimally to ensure that its endpoints do not belong to $\mathsf{FCVD}^{\star}(S)$ (Step 1). Then (in Step 2) it performs a check that allows to eliminate some execution paths, when it is guaranteed that e does not intersect $\mathsf{FCVD}^{\star}(S)$. If the check is passed, the search continues as follows. It performs a find-change query on uv that returns a point w (Steps 3–4). If $out \in type(w)$, the algorithm stops (Step 6), since it is guaranteed that $e \cap \mathsf{FCVD}^{\star}(S) = \emptyset$ by Lemma 13b. Otherwise, the algorithm proceeds. In case w is in $\mathsf{FCVD}^{\star}(S)$, the face containing w is traced (Step 8). In both cases (w is in $\mathsf{FCVD}^{\star}(S)$ or not), the algorithm calls itself recursively for two disjoint subsegments uq and $q'v$ of e (Steps 12–13). Initially, the algorithm is called for $e = uv$.

We remark that the faces of $\mathsf{FCVD}^{\star}(S)$ intersecting e are not found in a left-to-right or any other "natural" order, as the find-change query finds *some* point of the desired property. This is the reason why, after finding a changing point $w \in e$, the algorithm continues searching on both sides of w. We remark as well that, at every recursive call of the algorithm, the function $type(\cdot)$ is re-defined according to the extremes of the segment on which the procedure is called. Therefore, for a point $x \in e$, the value of $type(x)$ may change as the algorithm proceeds. Below we will refer to it as "$type(x)$ with respect to e", except for the cases when there is no ambiguity.

4.3 Correctness

We now prove that the algorithm presented in Sections 4.1 and 4.2 is correct. First we show that Find-change query is well-defined:

Lemma 14. *If the pair (t, s) satisfies the conditions of the input of Find-change query, then there exists a point w in the segment ts such that w is a changing point, $in \in type(w)$, or $out \in type(w)$.*

Proof. Suppose for the sake of contradiction that there is no such point w. Recall that there are five generic types for the points of e, namely $\{in, out, \tilde{l}, \tilde{r}, mm\}$. Since the lemma is not fulfilled, segment ts can be partitioned into open subsegments of points whose only type is either $\tilde{r}$ or $\tilde{l}$, and two consecutive such subsegments are separated by a point, whose type contains mm (the leftmost and the rightmost subsegment includes one border point t and s respectively, so they are half-open). Therefore, when traveling along ts from left to right, we encounter two consecutive subsegments of ts such that the only type of points in the first one is $\tilde{r}$, and the only type in the second one is $\tilde{\ell}$. Let w' be the point that separates these two subsegments; $mm \in type(w')$. We next argue that in this situation w' is a changing point, which will yield a contradiction.

Suppose that mm with multiplicity one is the only element in $type(w')$ (due to our assumption that no four endpoints of segments in S are cocircular, mm cannot be in $type(w')$ with multiplicity greater than one). Let cc' be the segment in S such that $mm \in type(w')$ caused by cc'. We assume without loss of generality that c' is contained in $D_h(u) \setminus D_h(v)$, and c in $D_h(v) \setminus D_h(u)$. Then $w' \in \mathsf{fcreg}(cc')$ and the farthest-color disk $D_f(w')$ has on the boundary c, c', and no other endpoint of any segment in S. Consequently, $w'_\ell \in \mathsf{fcreg}(c')$. This implies that $type(w'_\ell) = \{ml\}$, which contradicts the fact that the subsegment to the left of w' has only type $\tilde{r}$. Thus, there is another element in $type(w')$ apart from mm, and this element can only be $\tilde{r}$. Arguing analogously with w'_r, we find that $\tilde{l} \in type(w')$. We conclude that w' is a changing point, and arrive to a contradiction. $\square$

The following lemma refers to Step 2 of the algorithm SEARCH IN e:

Lemma 15. *Let $e = uv$, where $u, v \notin \mathsf{FCVD}^\star(S)$. If there is an element x in $type(u)$ such that $x \neq \tilde{r}$ or there is an element y in $type(v)$ such that $y \neq \tilde{l}$, then $uv \cap \mathsf{FCVD}^\star(S) = \emptyset$.*[5]

Proof. Recall that, if a point w belongs to $\mathsf{FCVD}^\star(S)$, then $D_h(w)$ contains at least one endpoint of every segment in S. Since $u \notin \mathsf{FCVD}^\star(S)$, each segment $cc' \in S$ such that $u \in \overline{\mathsf{fcreg}}(cc')$ has both endpoints outside $D_h(u)$. Therefore $in \notin type(u)$, $\tilde{l} \notin type(u)$, and $mm \notin type(u)$. Analogously, since $v \notin \mathsf{FCVD}^\star(S)$, we have that $in \notin type(v)$, $\tilde{r} \notin type(v)$, and $mm \notin type(v)$. On the other hand, by Lemma 13b if out belongs to $type(u)$ or to $type(v)$, then $uv \cap \mathsf{FCVD}^\star(S) = \emptyset$. $\square$

The following lemma proves that the algorithm SEARCH IN e is correct:

Lemma 16. *The algorithm* SEARCH IN e *computes all faces of* $\mathsf{FCVD}^\star(S)$ *intersected by* e.

Proof. Let $e = uv$. Segment uv satisfies the input condition of the algorithm: If uv is shrunk infinitesimally from both sides, its endpoints are outside $\mathsf{FCVD}^\star(S)$. This follows from Steps 13–14 of the algorithm COMPUTING $\mathsf{FCVD}^\star(S)$.

The condition checked in Step 2 of the algorithm is justified by Lemma 15: If the condition is not satisfied, then $uv \cap \mathsf{FCVD}^\star(S) = \emptyset$ and we can safely stop the search at Step 15. Otherwise, the input conditions of the find-change query are satisfied.

By Lemma 14, the find-change query is well-defined, and there are three possible outputs when we perform it on (u, v). If the query returns a point w such that $out \in type(w)$, then $e \cap \mathsf{FCVD}^\star(S) = \emptyset$ due to Lemma 13b. In this case, the algorithm stops (see Steps 5–6). Otherwise, a face of $\mathsf{FCVD}^\star(S)$ containing w is traced if and only if $w \in \mathsf{FCVD}^\star(S)$.

The algorithm is called recursively for two subsegments of uv, which are uq and $q'v$ (Steps 12–13). We now show that each of these subsegments satisfies the input condition of the algorithm. Indeed, points u_r and v_ℓ are outside $\mathsf{FCVD}^\star(S)$ because uv was satisfying this condition in the first place. Points q and q' are determined in Step 9 or Step 11, depending on the condition in Step 7. In particular, if the condition is satisfied, i.e., w lies in a face f of $\mathsf{FCVD}^\star(S)$, then $e \cap f$ is a line segment (see Lemma 11). In this case, Step 9 is executed, and it assigns to q and q' the left and right endpoints of the segment $e \cap f$, respectively. This implies that q_ℓ and q'_r are outside $\mathsf{FCVD}^\star(S)$. Otherwise, Step 11 assigns both q, q' to w and, since w is not in $\mathsf{FCVD}^\star(S)$, neither are w_ℓ, w_r.

In Theorem 3, we analyze the running time of the algorithm, and in particular we prove that the algorithm terminates. $\square$

[5] Note that it is not always possible to shrink uv infinitesimally so that $type(u)$ and $type(v)$ consist of one element each. In particular, this is not possible when uv lies on an edge of $\mathsf{FCVD}(S)$, and this situation does not contradict our general position assumption.

4.4 Running time

Below we analyze the time complexity of the algorithm to compute $\mathsf{FCVD}^{\star}(S)$, and thus the time complexity to solve the stabbing circle problem.

We start with the following result concerning the find-change query.

Lemma 17. *A find-change query can be performed in $O(\log^2 n)$ time.*

Proof. For a pair (t, s) of points, we perform the find-change query as follows. We use a point-location data structure for $\mathsf{FCVD}(S)$, such that the point location for a query point q is performed by a sequence of $O(\log n)$ atomic questions of the form "is q above or below (respectively, to the left or right of) a line ℓ?" (e.g., the data structure by Edelsbrunner et al. [17], or the one by Kirkpatrick [20]). Notice that, in our case, instead of a fixed point q, we only have a pair (t, s) such that the segment ts contains a sought point (a changing point, a point whose type contains *in* or a point whose type contains *out*). An atomic question is processed as follows. If $ts \cap \ell = \emptyset$, the answer is the same for any point in ts, and we continue with the pair (t, s). Otherwise, let point p be $ts \cap \ell$. First, the standard point location for p in $\mathsf{FCVD}(S)$ gives us $type(p)$. If $\tilde{r} \in type(p)$ and $\tilde{l} \notin type(p)$, we continue with the pair (p, s). Symmetrically, if $\tilde{l} \in type(p)$ and $\tilde{r} \notin type(p)$, we continue with the pair (t, p). If $type(p) = \{mm\}$, then we continue with either (p, s) or (t, p), because they both satisfy the input condition of the query. Otherwise, we stop the procedure, and return p. Clearly, this happens in one of the following cases: (i) $\{\tilde{l}, \tilde{r}\} \subseteq type(p)$; (ii) $in \in type(p)$; or (iii) $out \in type(p)$.

Answering one atomic question within the procedure takes $O(\log n)$ time, and the whole find-change query takes $O(\log^2 n)$ time.
A similar idea of simulating a point location for an unknown point is used in Cheong et al. [9], and in Cheilaris et al. [8, Section 7]. □

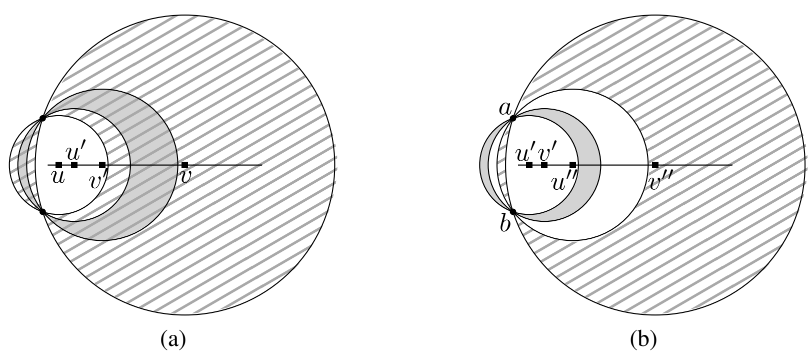

Fig. 16. Illustration for the proof of Lemma 18.

Lemma 18. *Let e be a segment on a pure edge of $\mathsf{HVD}(S)$, and cc' be a segment in S.*

(a) If cc' is of type middle for a segment $e' \subset e$, then cc' is of type middle for e.

(b) Let e', e'' be two disjoint subsegments of e. Then cc' is of type middle for at most one of e', e''.

Proof. Let $e = uv$, $e' = u'v'$, and $e'' = u''v''$.

(a) If cc' is of type middle for e', then one of its endpoint is in $D_h(u') \setminus D_h(v')$ and the other one is in $D_h(v') \setminus D_h(u')$ (see the two gray areas in Figure 16a). These areas are contained in $D_h(u) \setminus D_h(v)$ and $D_h(v) \setminus D_h(u)$, respectively (shown with tiling pattern in Figure 16a).

(b) Suppose cc' is of type middle for both e' and e''. One of the endpoints of cc' lies in $D_h(u') \setminus D_h(v')$ and the other one in $D_h(v') \setminus D_h(u')$. Analogously, one of its endpoints lies in $D_h(u'') \setminus D_h(v'')$ and the other one in $D_h(v'') \setminus D_h(u'')$. But these four areas are disjoint, see shaded and tiled areas in Figure 16b (note that the endpoints a, b of two segments in S such that $e \subset bis(a, b)$ belong to neither of these four areas since all the disks $D_h(\cdot)$ are closed). We arrive to a contradiction.

□

We now prove that the running time of SEARCH IN e is related to the number of segments of S of type middle for e, denoted by m_e.

Lemma 19. *Let e be a segment on a pure edge of $\mathsf{HVD}(S)$, such that the algorithm* SEARCH IN *is called for e. Then number of recursive calls performed by the algorithm* SEARCH IN *e is $O(1 + m_e)$.*

Proof. Let $e = uv$; assume without loss of generality that e is horizontal and u is its left endpoint. Consider the recursion tree of the algorithm SEARCH IN e. Observe that this tree is a proper binary tree (each call of the algorithm causes either zero or two recursive calls). Thus the total number of its nodes is linear in the number of its intermediate nodes. Therefore, it remains to prove that the number of intermediate nodes of the recursion tree of SEARCH IN e is $O(1 + m_e)$.

We will charge each intermediate node of the recursion tree to an endpoint of some segment of type middle for e, and each such endpoint will be charged at most once. Our charging scheme is symmetric: the nodes that are left children of their parents are charged to right endpoints, and vice versa. Thus we will discuss the charging scheme for the former case (left children) only. We base on the fact that for each segment of type *middle* for e there is at most one moment during the course of the algorithm when this segment becomes of type *left*. Formally, for any segment $cc' \in S$ that is of type *middle* for e, there is at most one node N in the recursion tree of the algorithm, such that for the subsegment of e that corresponds to N, cc' is of type *middle*, and for the subsegment of e that corresponds to the left child of N, cc' is of type *left*. This is easy to see from the following simple observations. First, if cc' has stopped being middle for the left child of N, then cc' stops being middle for the right child of N as well (it becomes of type *right*). Further, if there were two such nodes N and to N' for which the above property would hold, then N and N' would belong to one root-to-leaf path in the recursion tree (otherwise, the subsegments of e corresponding to N and N' would be disjoint, and due to Lemma 18b cc' could not be of type middle for both of them). Suppose that N is an ancestor of N'. On one root-to-leaf path, the segment on which the algorithm is called only shrinks and never expands, and thus by Lemma 18a after cc' has stopped being middle it never can be middle again (in particular, it cannot be middle for the segment corresponding to N'). We charge the left child of such node N to the right endpoint of cc'. We proceed with the technical details.

Consider a non-leaf node N of the recursion tree; let $u'v'$ be the subsegment of e that corresponds to node N. Since N is not a leaf, the algorithm has found a point $w \in u'v'$ such that either (1) $w \notin \mathsf{FCVD}^{\star}(S)$, w is a changing point, and *out* $\notin type(w)$, or (2) $w \in \mathsf{FCVD}^{\star}(S)$. We consider these two cases separately.

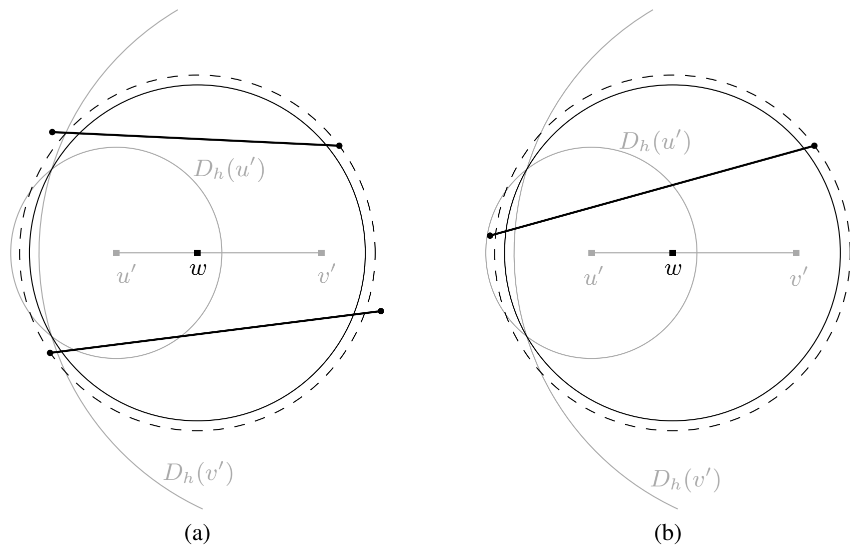

Fig. 17. Illustration for the proof of Lemma 19: the case when $w \notin \mathsf{FCVD}^{\star}(S)$ (case 1). Disks $D_h(u')$, $D_h(v')$ (gray lines), $D_h(w)$ (black solid lines); $D_f(w)$ (black dashed lines). (a) *right* $\in type(w)$ caused by cc': two possibilities for the segment cc' (bold line segment): $w \in \mathsf{fcreg}(c')$ (above), and $w \in \mathsf{fcreg}(c)$ (below). (b) *mr* $\in type(w)$ caused by cc': a placement of cc' (bold line segment).

Case 1: $w \notin \mathsf{FCVD}^{\star}(S)$, w is a changing point, and out $\notin type(w)$ *(see Figure 17).* The two children of the node N are the recursive calls for $u'w$ and for wv' respectively, and we will consider only the former one (the left child of N). Let cc' be the segment in S that causes type $\tilde{r}$ for w with respect to $u'v'$; and let c' be the endpoint of cc' such that $c' \in D_h(v') \setminus D_h(u')$. There are two possibilities: cc' causes type *right*, or it causes type *mr* (both types are defined with respect to $u'v'$); see respectively Figures 17a and 17b. In the former case, the call of the algorithm for $u'w$ corresponds to a leaf of the recursion tree: indeed, *out* $\in type(w_\ell)$ with respect to $u'_r w_\ell$ caused by cc' (see Figure 17a), and the algorithm will stop at Line 2. Thus no charging is needed for the left child of node N in this case. Figure 17a shows two possible placements of cc', for both of which the argument holds.

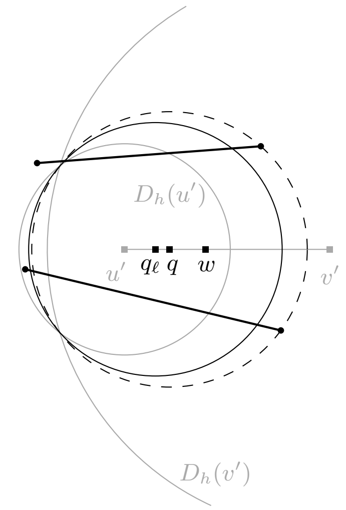

Fig. 18. Illustration for the proof of Lemma 19, the case when $w \in \mathsf{FCVD}^\star(S)$ (case 2). Disks $D_h(u')$, $D_h(v')$ (gray lines), $D_h(q) = D_f(q)$ (black dashed lines); $D_h(q_\ell)$ (black solid lines). Two possibilities for the segment cc' (bold line segment).

Consider the latter case, i.e., when cc' causes type *mr* with respect to $u'v'$. Observe that this may only happen when $w \in \mathsf{fcreg}(c')$: otherwise, cc' would by definition cause type *ml* or *mm*, rather than *mr*. This implies that cc' is of type *left* for $u'_r w_\ell$, see Figure 17b. We charge the call of the algorithm for segment $u'w$ (i.e., the left child of node N) to the endpoint c' of cc'.

Case 2: $w \in \mathsf{FCVD}^\star(S)$ (see Figure 18). The left child of node N in the recursion tree is the call of the algorithm for subsegment $u'q$ of $u'v'$, where q is the first point of the boundary of $\mathsf{FCVD}^\star(S)$ that is encountered when traversing e from w in the left direction. Since $q \in \mathsf{FCVD}^\star(S)$, each segment in S has at least one endpoint inside $D_h(q)$. Let cc' be a segment in S such that $q_\ell \in \overline{\mathsf{fcreg}}(cc')$; specifically let c' be the endpoint of cc' such that $q_\ell \in \overline{\mathsf{fcreg}}(c')$. Since q belongs to the boundary of $\mathsf{FCVD}^\star(S)$, $q_\ell \notin \mathsf{FCVD}^\star(S)$, and thus both c and c' are outside $D_h(q_\ell)$. This implies that c' lies in the portion of the boundary of $D_h(q)$ contained in $D_h(v')$, and c lies outside $D_h(q)$ (see Figure 18). Now, there are again two possibilities: cc' is of type *right* for $u'v'$, or cc' is of type *middle* for $u'v'$. In the former case, cc' is of type *out* for $u'_r q_\ell$, thus the call of the algorithm on $u'q$ is a leaf in the recursion tree, and no charging is needed for the left child of node N. Otherwise, i.e. if cc' is of type *middle* for $u'v'$, cc' is of type *left* for $u'_r q_\ell$, and we charge the call of the algorithm on $u'q$ (the left child of node N) to c'. $\square$

The next lemma allows to bound the total time of all executions of Step 11 of Computing $\mathsf{FCVD}^\star(S)$:

Lemma 20. *Let $f_1, f_2, \ldots, f_k$ be the faces of $\mathsf{FCVD}^\star(S)$. The total number of edges of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ intersected by these faces is $O(k + |\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$.*

Proof. Let I denote the total sum $\sum_{i=1}^{k} I(f_i)$, where $I(f_i)$ denotes the number of edges of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ that are intersected by f_i. We need to show that $I = O(k + |\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$. By Lemma 9, for each face f_i, $f_i \cap \mathsf{HVD}(S)$ is a connected component. Such connected component can be either (1) a portion of a single edge of $\mathsf{HVD}(S)$; or (2) a component containing portions of at least two edges. Depending on this, we call face f_i a type-1 face, or a type-2 face respectively. Let t be the number of type-1 faces, and r be the number of type-2 faces; then $t + r = k$. Clearly, all type-1 faces contribute t to the total sum I. Further, one edge of $\mathsf{HVD}(S)$ intersects at most two type-2 faces. Thus all type-2 faces contribute at most $2 * |\mathsf{HVD}(S)|$ to I. Since $t \leq k$, $\mathsf{HVD}(S)$ contributes at most $k + 2 * |\mathsf{HVD}(S)|$ to I. By an analogous argument, $\mathsf{FCVD}(S)$ contributes at most $k + 2 * |\mathsf{FCVD}(S)|$ to I. The claim follows. $\square$

We are finally ready to prove the main theorem of this section.

Let m denote the number of pairs formed by a segment $aa' \in S$ and a pure edge e of $\mathsf{HVD}(S)$ such that aa' is of type *middle* for e. That is, m is the total sum of m_e for all pure edges e of $\mathsf{HVD}(S)$. Let $\mathcal{T}_{\mathsf{HVD}(S)}$ and $\mathcal{T}_{\mathsf{FCVD}(S)}$ denote the time to compute $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$, respectively. Let $\mathcal{T}_q$ denote the time to answer the find-change query. In this paper, $\mathcal{T}_q$ is $O(\log^2 n)$.

Theorem 3. *Let S be a set of n segments in the plane in general position. Then, $\mathsf{FCVD}^\star(S)$ can be computed in time $O(\mathcal{T}_{\mathsf{HVD}(S)} + \mathcal{T}_{\mathsf{FCVD}(S)} + (|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + m)(\log n + \mathcal{T}_q))$.*

Proof. Computing $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ (Step 1 of COMPUTING $\mathsf{FCVD}^\star(S)$) requires time $\mathcal{T}_{\mathsf{HVD}(S)} + \mathcal{T}_{\mathsf{FCVD}(S)}$.

After computing the diagrams, in time $O(|\mathsf{HVD}(S)| \log n)$ and $O(|\mathsf{FCVD}(S)| \log n)$ we can preprocess them to answer point-location queries in $O(\log n)$ time [17,20]. Then Steps 3 and 9 of the algorithm require logarithmic time (see Section 4.1).

We also preprocess $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ to answer ray-shooting queries in $O(\log n)$ time. The preprocessing can be done in time $O(|\mathsf{HVD}(S)| \log n)$ and $O(|\mathsf{FCVD}(S)| \log n)$, respectively [7]. After such preprocessing, Step 11 requires $O((1+I(f)) \log n)$ time, where $I(f)$ is the number of edges of $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$ intersected by f (see Section 4.1 for more details on Step 11). By Lemma 20, the total time spent for Step 11 is $O((|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + |\mathsf{FCVD}^\star(S)|) \log n)$.

After a point on ∂f is known, ∂f can be traced in $\mathsf{HVD}(S)$ and $\mathsf{FCVD}(S)$. Thus the total time required for tracing the boundary of all the faces of $\mathsf{FCVD}^\star(S)$ (in Steps 6 and 11 of COMPUTING $\mathsf{FCVD}^\star(S)$, and Step 8 of SEARCH IN e) is $O((|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + |\mathsf{FCVD}^\star(S)|) \log n)$.

Therefore, the total time complexity of Steps 2–12 of COMPUTING $\mathsf{FCVD}^\star(S)$ is $O((|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + |\mathsf{FCVD}^\star(S)|) \log n)$.

The third loop (Steps 13–14) of COMPUTING $\mathsf{FCVD}^\star(S)$ performs $O(|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$ iterations. Indeed, the number of faces computed in Steps 2–12 is $O(|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)|)$, and each face has connected intersection with $\mathsf{HVD}(S)$ by Lemma 9.

In Step 14, SEARCH IN e is called. By Lemma 19, the total number of recursive calls caused by SEARCH IN $e = uv$ is $O(1 + m_e)$. Summing up m_e for all pieces e of edges of $\mathsf{HVD}(S)$ on which the procedure is called results in $O(m)$ due to Lemma 18b. Since any face of $\mathsf{FCVD}^{\star}(S)$ computed by SEARCH IN e intersects only one edge of $\mathsf{HVD}(S)$ and one edge of $\mathsf{FCVD}(S)$ and its boundary has constant complexity (see Lemma 11), Step 8 requires constant time. Therefore, one execution of SEARCH IN e, except for the recursive calls, is dominated by the find-change query, and thus requires $O(\log^2 n)$ time (see Lemma 17). Thus the total time required by SEARCH IN e is $O((1+m_e)\log^2 n)$. This globally amounts to $O((|\mathsf{HVD}(S)|+|\mathsf{FCVD}(S)|+m)\log^2 n)$ time for Steps 13–14 of COMPUTING $\mathsf{FCVD}^{\star}(S)$. More precisely, this time is $O((|\mathsf{HVD}(S)|+|\mathsf{FCVD}(S)|+m)(\mathcal{T}_q+\log n))$. $\square$

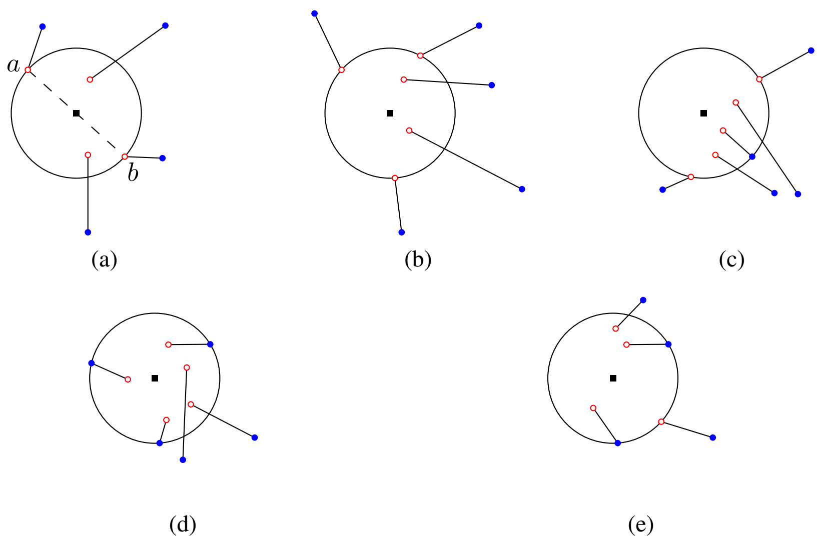

Fig. 19. All possibilities for the stabbing circles of minimum and maximum radius.

The next lemma shows how to find the largest and smallest stabbing circles, once $\mathsf{FCVD}^{\star}(S)$ is known. We observe that, in some cases, the stabbing circle of minimum radius does not exist, because any stabbing circle can be shrunk by decreasing its radius or slightly moving its center. Moreover, the "limit" circle is not stabbing because the closed disk induced by the circle contains both endpoints of a segment in S (see Case 1c of the proof of Lemma 21 below and Figure 19c). Similarly, it is easy to see that the stabbing circle of maximum radius never exists, since any stabbing circle can be slightly enlarged, but the "limit" circle is not stabbing (see Case 2 of the proof of Lemma 21 and Figure 19d,e). In these cases, even though these circles are not stabbing, to simplify the notation we call these "limit" circles the *stabbing circles with infimum* or *supremum radius*.

Lemma 21. *After computing $\mathsf{FCVD}^{\star}(S)$, the stabbing circles of minimum (or infimum) and supremum radius can be determined in time $O((|\mathsf{FCVD}^{\star}(S)|+|\mathsf{HVD}(S)|+|\mathsf{FCVD}(S)|)\log n)$.*

Proof. We list all the possibilities for such circles below; see Figure 19.

Given a segment aa' and a circle c, we call "red" the endpoint of aa' which is closer to the center of c, and we call "blue" the other endpoint (ties are broken arbitrarily). Notice that, if c is stabbing, then the red endpoint of aa' lies inside the closed disk induced by c, and the blue endpoint lies outside. Standard geometric arguments show the following.

1. The stabbing circle of minimum or infimum radius is a circle passing through:
   (a) Two red points that are diametrically opposite. In this case the center of the circle is at the intersection point between an edge of $\mathsf{FCVD}(S)$ separating two regions $\mathsf{fcreg}(a)$ and $\mathsf{fcreg}(b)$ (for some $aa', bb' \in S$), and the segment connecting a to b, see Figure 19a. This is a stabbing circle of minimum radius.
   (b) Three red points. The center is at a vertex of $\mathsf{FCVD}(S)$. This is also a stabbing circle of minimum radius.
   (c) Two red points and one blue point. The center is on the boundary of $\mathsf{FCVD}^\star(S)$. This is a stabbing circle of infimum radius.
2. The stabbing circle of supremum radius is a circle passing through:
   (a) Three blue points. In this case the center of the circle is at a vertex of $\mathsf{HVD}(S)$.
   (b) Two blue points and one red point. The center is on the boundary of $\mathsf{FCVD}^\star(S)$.

The stabbing circles with minimum (or infimum) and supremum radius can thus be found by checking the vertices of $\mathsf{HVD}(S)$ or $\mathsf{FCVD}(S)$ lying in $\mathsf{FCVD}^\star(S)$, the edges of $\mathsf{FCVD}(S)$ intersecting $\mathsf{FCVD}^\star(S)$, and the boundary of $\mathsf{FCVD}^\star(S)$. Since in all cases the radius of the corresponding circle can be computed in $O(\log n)$ time, the claim follows. $\square$

As a conclusion, we put together the above results and connect them with the stabbing circle problem as stated in Section 1. By Lemma 7, distinct faces of $\mathsf{FCVD}^\star(S)$ correspond to combinatorially different stabbing circles. Thus, the above results yield the following:

Corollary 3. *Let S be a set of n segments in the plane in general position. All the combinatorially different stabbing circles for S, and the ones with minimum and maximum radius, can be computed in $O(\mathcal{T}_{\mathsf{HVD}(S)} + \mathcal{T}_{\mathsf{FCVD}(S)} + (|\mathsf{HVD}(S)| + |\mathsf{FCVD}(S)| + m)\log^2 n)$ time.*

5 Parallel Segments

Let S be a set of parallel segments. The goal of this section is to prove the following theorem.

Theorem 4. *The stabbing circle problem for a set S of n parallel segments in general position can be solved in $O(n \log^2 n)$ time and $O(n)$ space.*

We remark that the $O(\log^2 n)$ factor in the time complexity of Theorem 4 is contributed by the find-change query. If this query can be answered in $O(\log n)$ time then the stabbing circle problem for n parallel segments can be solved in $O(n \log n)$ time.

We prove Theorem 4 using Theorem 3 and Corollary 3. Since the segments in S are parallel, they must by pairwise disjoint, and thus $\mathsf{HVD}(S)$ is an instance of *abstract Voronoi diagrams*; hence, $|\mathsf{HVD}(S)|$ is $O(n)$ and $\mathcal{T}_{\mathsf{HVD}(S)}$ is $O(n \log n)$ [21]. In Section 5.1 we show that $|\mathsf{FCVD}(S)|$ is also $O(n)$ and $\mathcal{T}_{\mathsf{FCVD}(S)}$ is $O(n \log n)$. In Section 5.2 we show that $m = O(n)$. Thus, the algorithm of Section 4 for this particular case has time complexity $O(n(\log n + \mathcal{T}_q))$, and since $\mathcal{T}_q$ is $O(\log^2 n)$ (see Lemma 17) we derive Theorem 4.

5.1 The farthest-color Voronoi diagram for a set of parallel segments

Lemma 22. *If all segments in S are parallel, then for each $aa' \in S$, $bis(a, a')$ contributes at most one internal edge to* $\mathsf{FCVD}(S)$.

Proof. Suppose for the sake of contradiction that there are two internal edges e and e' in $\mathsf{FCVD}(S)$ which are portions of $bis(a, a')$. Assume without loss of generality that the segments in S are vertical, thus, $bis(a, a')$ is horizontal, and e is to the left of e'. Let $v, v' \in bis(a, a')$ be the mixed vertices that are respectively the right endpoint of e and the left endpoint of e'. The farthest-color disk $D_f(v)$ contains an endpoint of every segments in S, and $\partial D_f(v)$ passes through points a, a' and an endpoint b of some segment $bb' \in S$. Observe that bb' is to the left of aa', and $bb' \cap D_f(v) = \{b\}$ (see Figure 20 (left), where $\mathsf{fcreg}(aa')$ is shown shaded, and $bis(a, a')$ is shown dashed). Since v' is to the right of v, the closed disk centered at v' with radius $d(v', a)$ does not contain b nor b', thus, $v' \notin \overline{\mathsf{fcreg}}(aa')$. We obtain a contradiction. □

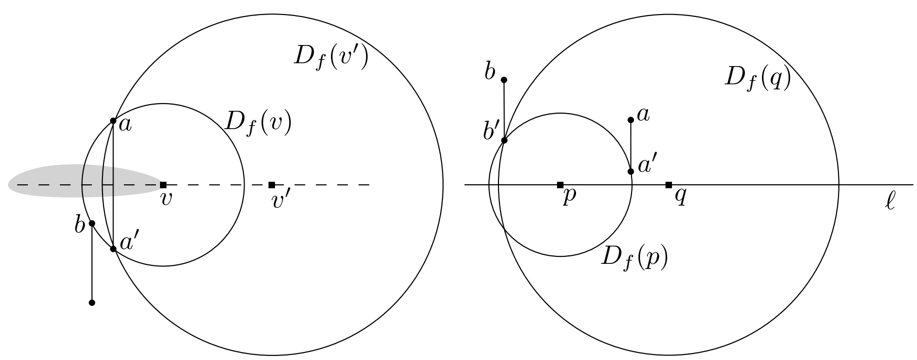

Fig. 20. Left: Illustration for the proof of Lemma 22. Right: Illustration for the proof of Lemma 24.

Lemma 23. *If all segments in S are parallel, the structural complexity of* $\mathsf{FCVD}(S)$ *is* $O(n)$.

Proof. By Lemma 5, the number of unbounded faces of $\mathsf{FCVD}(S)$ is $O(n)$. Since all segments in S are parallel, by Lemma 22, the number of all internal edges in $\mathsf{FCVD}(S)$ is at most n. By Lemma 3, every bounded face of $\mathsf{FCVD}(S)$ has an internal edge on its boundary, thus there are at most $2n$ bounded faces in $\mathsf{FCVD}(S)$. Thus, $\mathsf{FCVD}(S)$ is a planar graph with $O(n)$ faces whose vertices have degree three. By Euler's formula, $|\mathsf{FCVD}(S)| = O(n)$. ☐

Lemma 24. *If all segments in S are parallel, then $\mathsf{FCVD}(S)$ can be computed in $O(n \log n)$ time.*

Proof. Assume without loss of generality that the segments in S are vertical. We use the divide-and-conquer technique: We divide S by a vertical line into two halves, S_{left} and S_{right}, and recursively compute $\mathsf{FCVD}(S_{left})$ and $\mathsf{FCVD}(S_{right})$. Below we prove that the merge curve between $\mathsf{FCVD}(S_{left})$ and $\mathsf{FCVD}(S_{right})$ is y-monotone. Such a y-monotone merge curve can be constructed in $O(n)$ time by standard arguments on the divide-and-conquer construction of Voronoi diagrams, see e.g., [5]. The claim follows.

Suppose for the sake of contradiction, that the merge curve is not y-monotone, that is, there are two points p, q on the merge curve with the same y-coordinate. Since p and q lie on the merge curve, the farthest-color radius of p with respect to S_{left} equals the farthest-color radius of p with respect to S_{right}, and the same holds for q. Let ℓ be the horizontal line through p and q. We redefine p and q as the left-most and second left-most points on ℓ that lie on the merge curve, respectively. In $\mathsf{FCVD}(S)$, the point at minus infinity on ℓ lies in the farthest-color region of a segment from S_{right}. Thus p lies on the boundary between $\mathsf{fcreg}(aa')$ (to the left of p) and $\mathsf{fcreg}(bb')$ (to the right of p) such that $aa' \in S_{right}$ and $bb' \in S_{left}$; see Figure 20 (right). Then the farthest-color disk $D_f(q)$ intersects or touches bb', and touches a segment $cc' \in S_{right}$. But in this case cc' is outside the farthest-color disk of p. We obtain a contradiction. ☐

5.2 Segments of type middle for a set of parallel segments

The key result of this subsection is the following:

Lemma 25. *A segment $gg' \in S$ is of type middle for at most one pure edge of $\mathsf{HVD}(S)$.*

We first prove an easy property of the segments of type middle:

Lemma 26. *Suppose that all segments in S are vertical. Let e be a pure edge of $\mathsf{HVD}(S)$ in the boundary of $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$, for two segments $aa', bb' \in S$. Suppose that segment $gg' \in S$ is of type middle for e. Then:*

(a) $x(a) \neq x(b)$.
(b) $\min\{x(a), x(b)\} < x(g) < \max\{x(a), x(b)\}$.

Proof. To prove (a), we suppose for the sake of contradiction that $x(a) = x(b)$. Then e is horizontal. Let u and respectively v be the left and right endpoints of e. Since the segment gg' is of type middle, it has one endpoint to the left of the line $x = x(a)$, in $D_h(u) \setminus D_h(v)$, and the other endpoint to the right of the line $x = x(a)$, in $D_h(v) \setminus D_h(u)$ (see Figure 21, left). This contradicts the fact that gg' is vertical.

To prove (b), we suppose without loss of generality that $x(a) < x(b)$. Then e is not horizontal, and we denote by u and respectively v the top and bottom endpoints of e.

For any disk D, we can divide its boundary ∂D into the left-most point of ∂D, the open circular arc containing the upper half portion of ∂D (called *top chain*), the right-most point of ∂D, and the open circular arc containing the lower half portion of ∂D (called *bottom chain*). Since aa' is vertical and $D_h(u)$ contains both a and a', a is not the left-most or right-most point of $\partial D_h(u)$. Suppose that a belongs to the top chain of $\partial D_h(u)$. Then we deduce that $y(a') < y(a)$. Since $D_h(v)$ also contains both a and a', we get that a belongs to the top chain of $\partial D_h(v)$. So a belongs to the top (resp., bottom) chain of $\partial D_h(u)$ if and only if a belongs to the top (resp., bottom) chain of $\partial D_h(v)$. The same argument applies to b.

Now there are several possibilities, depending on a and b being in the top or bottom chains of $\partial D_h(u)$ and $\partial D_h(v)$. The arguments for all cases are similar, so we only explain the case where a and b belong to the top chains of $\partial D_h(u)$ and $\partial D_h(v)$. In this case, u and v lie in the portion of $bis(a, b)$ below the lines $y = y(a)$ and $y = y(b)$ (and recall that $y(u) > y(v)$) (see Figure 21, center). Either g or g' lies in $D_h(u) \setminus D_h(v)$, so in particular in the portion of $D_h(u)$ above the segment ab. Since a and b belong to the top chain of $\partial D_h(u)$, this portion lies between the lines $x = x(a)$ and $x = x(b)$. Thus we obtain $x(a) < x(g) < x(b)$. $\square$

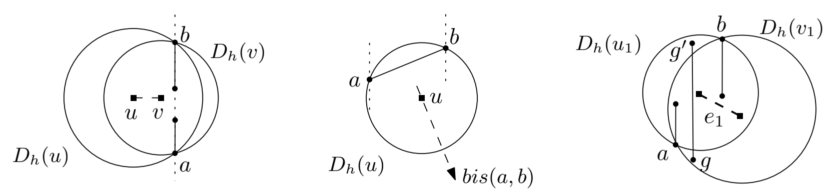

Fig. 21. Left: Case where $x(a) = x(b)$. Center: Either g or g' lies in a portion of $D_h(u)$ between the lines $x = x(a)$ and $x = x(b)$. Right: The segment gg' is of type middle for e_1.

We are now ready to prove Lemma 25.

Proof (of Lemma 25). We assume that all segments in S are vertical. We proceed by contradiction. So let us assume that the segment gg' is of type middle for two pure edges of $\mathsf{HVD}(S)$, namely e_1 and e_2. Let e_1 be in the boundary of $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$, for two segments $aa', bb' \in S$ (see Figure 21, right). Analogously, e_2 is in the boundary of $\mathsf{hreg}(c)$ and $\mathsf{hreg}(d)$, for two segments $cc', dd' \in S$. By Lemma 26a, $x(a) \neq x(b)$ and $x(c) \neq x(d)$. Without loss of generality, we suppose that $x(a) < x(b)$ and $x(c) < x(d)$. We also assume that $y(g) < y(g')$.

Consider the disk having a, b and g on the boundary; this disk corresponds to a disk $D_h(w_1)$, for some point w_1 on the edge e_1. Analogously, the disk having c, d and g on the boundary corresponds to a disk $D_h(w_2)$, for some point w_2 on e_2.

Since, by Lemma 26b, $x(a) < x(g) < x(b)$, the line $x = x(g)$ intersects $\partial D_h(w_1)$ twice. One of these intersection points is g, and we next show that the second intersection point, called t_1, is above g. Since the line through a and b leaves g and g' on opposite sides, and since $x(a) < x(g) < x(b)$, the segment ab intersects the segment gg'. The intersection point lies above g and, by convexity, it is contained in $D_h(w_1)$. This implies that g is in the bottom chain of $\partial D_h(w_1)$ and, consequently, t_1 is above g (see Figure 22, left). Analogously, the second intersection point t_2 between $\partial D_h(w_2)$ and $x = x(g)$ also lies above g. Without loss of generality, we assume that $y(t_1) \geq y(t_2)$. We divide the rest of the argument into several cases.

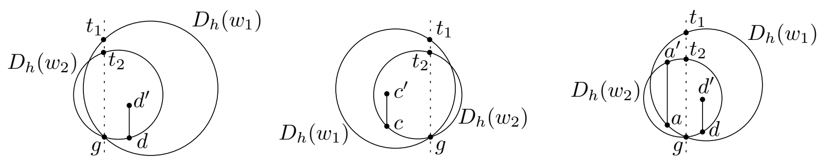

Fig. 22. Left: Case where the four segments are distinct and the second intersection point between $\partial D_h(w_1)$ and $\partial D_h(w_2)$ is to the left of $x = x(g)$. Middle: Case where the four segments are distinct and the second intersection point between $\partial D_h(w_1)$ and $\partial D_h(w_2)$ is to the right of $x = x(g)$. Right: Case where $a = c'$.

We start by considering the case where the four segments aa', bb', cc', and dd' are distinct. Because w_1 is in the boundary of $\mathsf{hreg}(a)$ and $\mathsf{hreg}(b)$ and the four segments are distinct, $D_h(w_1)$ contains at most one of $\{c, c'\}$ and at most one of $\{d, d'\}$. Analogously, $D_h(w_2)$ contains at most one of $\{a, a'\}$ and at most one of $\{b, b'\}$. Consequently, none of $D_h(w_1), D_h(w_2)$ contains the other, and $\partial D_h(w_1)$ and $\partial D_h(w_2)$ intersect at g and at a second point. If this second point lies to the left of the line $x = x(g)$, then the portion of $D_h(w_2)$ to the right of $x = x(g)$ is contained in $D_h(w_1)$ (see Figure 22, left). Consequently, dd' is in $D_h(w_1)$, yielding a contradiction. If the second intersection point between $\partial D_h(w_1)$ and $\partial D_h(w_2)$ lies to the right of the line $x = x(g)$, then cc' is in $D_h(w_1)$ (see Figure 22, center). If the intersection point lies on $x = x(g)$, then $y(t_1) = y(t_2)$, and we obtain that $D_h(w_1)$ contains cc' or dd'.

Let us look at the remaining cases. Since $x(a) < x(g) < x(b)$ and $x(c) < x(g) < x(d)$, the intervals $(x(a), x(b))$ and $(x(c), x(d))$ have non-empty intersection. This implies that $aa' \neq dd'$ and $bb' \neq cc'$ (and obviously $aa' \neq bb'$ and $cc' \neq dd'$). Therefore, the two remaining cases are $aa' = cc'$ and $bb' = dd'$. We divide the first one into two subcases, namely, $a = c$ and $a = c'$.

If $a = c$, the second intersection point between $\partial D_h(w_1)$ and $\partial D_h(w_2)$ is a, which lies to the left of the line $x = x(g)$. If $bb' \neq dd'$, then dd' is contained in $D_h(w_1)$, yielding a contradiction. If $bb' = dd'$, then, since b lies on $\partial D_h(w_1)$ and d lies on $\partial D_h(w_2)$, we have that $b \neq d$. But then $b = d'$ lies on the portion of $\partial D_h(w_1)$ to the right of $x = x(g)$, that is, outside $D_h(w_2)$, yielding a contradiction.

If $a = c'$, due to the assumption that $y(t_1) \geq y(t_2)$, we have that a is in the bottom chain of $\partial D_h(w_1)$ and $a' = c$ is in the top chain of $\partial D_h(w_2)$ (see Figure 22, right). Then the second intersection point between $\partial D_h(w_1)$ and $\partial D_h(w_2)$ lies to the left of

$x = x(a)$, and we also have that dd' is in $D_h(w_1)$ (if $bb' \neq dd'$) or that d' is outside $D_h(w_2)$ (if $bb' = dd'$).

In the last case, $bb' = dd'$. This case is symmetric to the previous one, and it also yields a contradiction. □

6 Lower Bound

Below we prove a lower bound for computing a stabbing circle of minimum radius for sets of segments which are parallel and have equal length.

Theorem 5. *The problem of computing a stabbing circle of minimum radius for a set of n parallel segments of equal length has an $\Omega(n \log n)$ lower bound in the algebraic decision tree model.*

Proof. The reduction, very similar to that of Theorem 6 in [3], is from $\mathsf{MAXGAP}(X)$. In our version, the input X consists of a set of n integers $x_1, \ldots, x_n$, and $\mathsf{MAXGAP}(X)$ is the problem of finding the maximum difference between consecutive elements of X.

Without loss of generality, we may assume $\min X = 1$. Let $x'_1 < x'_2 < \cdots < x'_n$ be the sorting of the elements of X. Then $x'_1 = 1$, and let $M = x'_n$. We construct a set S of parallel segments of equal length as follows: For every $x_i \in X$, we add a segment connecting point $(x_i, 0)$ to $(-(M+1) + x_i, 0)$. Additionally, we add two segments aa' and bb' such that $a = (-1/2, 0)$, $a' = (-(M+1) - 1/2, 0)$, $b = (1/2, 0)$, and $b' = ((M+1) + 1/2, 0)$.

Any stabbing circle for S of minimum radius contains a, b in its interior. Thus the possibilities for such a stabbing circle are: If the associated disk contains $a, b, (x'_1, 0), \ldots, (x'_n, 0)$, or $(-(M+1) + x'_1, 0), \ldots, (-(M+1) + x'_n, 0), a, b$, then it has diameter $M + 1/2$. If it contains $(-(M+1) + x'_{i+1}, 0), \ldots, (-(M+1) + x'_n, 0), a, b, (x'_1, 0), \ldots, (x'_i, 0)$ for $i < n$, then it has diameter $M + 1 - (x'_{i+1} - x_i)$. Since $\mathsf{MAXGAP}(X) \geq 1$, the stabbing circles of minimum radius belong to the last family. Thus $\mathsf{MAXGAP}(X)$ is equivalent to finding the stabbing circle for S of minimum radius.

The set S does not satisfy all the assumptions of this paper, since all endpoints are collinear. We construct a set S' obtained from S by translating every segment vertically by distinct values of at most $\varepsilon = 1/10$. Since ε is small compared to the difference between distinct values of diameters of different stabbing circles for S (which is at least $1/2$), a minimum stabbing circle for S' corresponds to a minimum stabbing circle for S which is combinatorially "the same". This proves that the lower bound also holds for the more restricted sets of segments considered in this paper. □

7 Conclusions and Future Work

The connection between the stabbing circle problem and the cluster Voronoi diagrams allows to solve the stabbing circle problem in an efficient way under certain conditions on S. In fact, this connection is a good tool for studying the stabbing circle problem for segment sets of particular types. So our goal is "to understand" the time and space complexities depending on the kind of segment sets we work with. In this paper, we have shown that our method allows to solve the stabbing circle problem in $o(n^2)$ time (that is, faster than in the general case) when all segments are parallel. The further open question is to investigate other segment sets for which this is also the case.

A very recent short abstract [13] is a preliminary step in this direction. We consider a set S of n disjoint segments that correspond to edges of the Delaunay triangulation of some fixed point set, and we show that the stabbing circle problem for S can be solved in $O(n \log n)$ time and $O(n)$ space in two cases: (i) all segments in S are parallel; (ii) all segments in S have equal length. It remains as an open problem whether the stabbing circle problem for such a set S (not necessarily satisfying (i) or (ii)) can be solved in subquadratic time. Another case for which it might be possible to achieve near-linear time is that of a set of pairwise-disjoint segments which are vertical or horizontal, or which are aligned to a given (constant) set of directions.

In a different vein, it would be interesting to find a way to charge the time and space complexities required to solve the stabbing circle problem directly to the complexity of $\mathsf{FCVD}^*(S)$, rather than to the number of segments of type middle for S.

Further, the find-change query in this paper is implemented in a straight forward way. It is an open problem whether this query could be handled more efficiently, in say $O(\log n)$ time. If this question was answered in the affirmative, the time complexity of our algorithms would improve by a logarithmic factor.

Finally, deriving a lower bound for the decision version of the stabbing circle problem remains open.

Acknowledgments. M. C. and C. S. were supported by projects MTM2015-63791-R (MINECO/FEDER) and Gen.Cat. DGR2014SGR46. E. K. and E. P. were supported by projects SNF 20GG21-134355, under the ESF EUROCORES program EuroGIGA/VORONOI, and SNF 200021E-154387. M. S. was supported by project LO1506 of the Czech Ministry of Education, Youth and Sports, and by project NEXLIZ CZ.1.07/2.3.00/30.0038, co-financed by the European Social Fund and the state budget of the Czech Republic.

References

1. Abellanas, M., Hurtado, F., Icking, C., Klein, R., Langetepe, E., Ma, L., Palop, B., Sacristán, V.: The farthest color Voronoi diagram and related problems. In: 17th European Workshop on Comput. Geom. (EuroCG'01). pp. 113–116 (2001), Tech. Rep. 002 2006, Univ. Bonn
2. Arkin, E.M., Dieckmann, C., Knauer, C., Mitchell, J.S., Polishchuk, V., Schlipf, L., Yang, S.: Convex transversals. Comput. Geom. 47(2), 224–239 (2014)
3. Arkin, E.M., Díaz-Báñez, J.M., Hurtado, F., Kumar, P., Mitchell, J.S.B., Palop, B., Pérez-Lantero, P., Saumell, M., Silveira, R.I.: Bichromatic 2-center of pairs of points. Comput. Geom. 48(2), 94–107 (2015)
4. Aurenhammer, F., Drysdale, R., Krasser, H.: Farthest line segment Voronoi diagrams. Inform. Process. Lett. 100, 220–225 (2006)
5. Aurenhammer, F., Klein, R., Lee, D.T.: Voronoi Diagrams and Delaunay Triangulations. World Scientific (2013)
6. Avis, D., Robert, J., Wenger, R.: Lower bounds for line stabbing. Inform. Process. Lett. 33(2), 59–62 (1989)

7. Chazelle, B., Edelsbrunner, H., Grigni, M., Guibas, L.J., Hershberger, J., Sharir, M., Snoeyink, J.: Ray shooting in polygons using geodesic triangulations. Algorithmica 12(1), 54–68 (1994)
8. Cheilaris, P., Khramtcova, E., Langerman, S., Papadopoulou, E.: A randomized incremental algorithm for the Hausdorff Voronoi diagram of non-crossing clusters. Algorithmica (2016), DOI 10.1007/s00453-016-0118-y
9. Cheong, O., Everett, H., Glisse, M., Gudmundsson, J., Hornus, S., Lazard, S., Lee, M., Na, H.: Farthest-polygon Voronoi diagrams. Comput. Geom. 44(4), 234–247 (2011)
10. Claverol, M.: Problemas geométricos en morfología computacional. Ph.D. thesis, Universitat Politècnica de Catalunya (2004)
11. Claverol, M., Garijo, D., Grima, C.I., Márquez, A., Seara, C.: Stabbers of line segments in the plane. Comput. Geom. 44(5), 303–318 (2011)
12. Claverol, M., Garijo, D., Korman, M., Seara, C., Silveira, R.I.: Stabbing segments with rectilinear objects. In: Kosowski, A., Walukiewicz, I. (eds.) FCT 2015. LNCS, vol. 9210, pp. 53–64. Springer (2015)
13. Claverol, M., Khramtcova, E., Papadopoulou, E., Saumell, M., Seara, C.: Stabbing circles for some sets of Delaunay segments. In: 32th European Workshop on Comput. Geom. (EuroCG'16). pp. 139–143 (2016)
14. Díaz-Báñez, J.M., Korman, M., Pérez-Lantero, P., Pilz, A., Seara, C., Silveira, R.I.: New results on stabbing segments with a polygon. Comput. Geom. 48(1), 14–29 (2015)
15. Edelsbrunner, H., Maurer, H., Preparata, F., Rosenberg, A., Welzl, E., Wood, D.: Stabbing line segments. BIT 22(3), 274–281 (1982)
16. Edelsbrunner, H., Guibas, L.J., Sharir, M.: The upper envelope of piecewise linear functions: algorithms and applications. Discrete Comput. Geom. 4, 311–336 (1989)
17. Edelsbrunner, H., Guibas, L.J., Stolfi, J.: Optimal point location in a monotone subdivision. SIAM J. Comput. 15(2), 317–340 (1986)
18. Edelsbrunner, H., Seidel, R.: Voronoi diagrams and arrangements. Discrete Comput. Geom. 1(1), 25–44 (1986)
19. Huttenlocher, D.P., Kedem, K., Sharir, M.: The upper envelope of Voronoi surfaces and its applications. Discrete Comput. Geom. 9(1), 267–291 (1993)
20. Kirkpatrick, D.: Optimal search in planar subdivisions. SIAM J. Comput. 12(1), 28–35 (1983)
21. Klein, R.: Concrete and Abstract Voronoi Diagrams, Lecture Notes in Computer Science, vol. 400. Springer-Verlag (1989)
22. Papadopoulou, E., Lee, D.T.: The Hausdorff Voronoi diagram of polygonal objects: a divide and conquer approach. Internat. J. Comput. Geom. Appl. 14(6), 421–452 (2004)
23. Papadopoulou, E.: The Hausdorff Voronoi diagram of point clusters in the plane. Algorithmica 40(2), 63–82 (2004)
24. Papadopoulou, E., Dey, S.K.: On the farthest line-segment Voronoi diagram. Internat. J. Comput. Geom. Appl. 23(06), 443–459 (2013)
25. Rappaport, D.: Minimum polygon transversals of line segments. Internat. J. Comput. Geom. Appl. 5(3), 243–256 (1995)