

Three Dimensional Euclidean Voronoi Diagrams of Lines with a Fixed Number of Orientations*

Vladlen Koltun[†] Micha Sharir[‡]

February 3, 2003

Abstract

We show that the combinatorial complexity of the Euclidean Voronoi diagram of n lines in \mathbb{R}^3 that have at most c distinct orientations is $O(c^3 n^{2+\varepsilon})$, for any $\varepsilon > 0$. This result is a step towards proving the long-standing conjecture that the Euclidean Voronoi diagram of lines in three dimensions has near-quadratic complexity. It provides the first natural instance in which this conjecture is shown to hold. In a broader context, our result adds a natural instance to the (rather small) pool of instances of general 3-dimensional Voronoi diagrams for which near-quadratic complexity bounds are known.

1 Introduction

Background. The Voronoi diagram of a set Γ of disjoint objects (‘sites’) in some space under some metric is a subdivision of the space into cells, one cell per site, such that the cell associated with a site $O \in \Gamma$ comprises the points in space for which O is closer (under the given metric) than all other sites of Γ .

The study of Voronoi diagrams in the plane has been very extensive over the past 20 years, and the structure of such diagrams is by now thoroughly understood. The study has covered diagrams for many kinds of sites, and for many kinds of metrics

*Work on this paper has been supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing). Work by Micha Sharir was also supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the U.S.-Israel Binational Science Foundation, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. This work is part of the first author’s Ph.D. dissertation, prepared under the supervision of the second author. A preliminary version of this paper is to appear in the Proceedings of the 18th ACM Symposium on Computational Geometry, 2002.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel. vladlen@tau.ac.il

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. sharir@cs.tau.ac.il

or distance functions, and has also considered other variants of the problem, such as k -th order diagrams, constrained Delaunay triangulations, and more. Surveys of the state of the art are given in Aurenhammer and Klein [4] and Fortune [10].

In contrast, Voronoi diagrams in three and higher dimensions have been much less studied, and many basic problems are still wide open. Most variants of planar Voronoi diagrams have linear complexity, which is usually a consequence of the planarity of the diagram. In three dimensions, a prevailing conjecture is that the complexity of Voronoi diagrams should be in general at most quadratic or near-quadratic in the number of sites. This is known to hold only for very few special cases, including the cases of point sites under the Euclidean metric [17, 22], of point sites under any ‘polyhedral’ metric or distance function (i.e., distance functions induced by a convex polytope with $O(1)$ facets; see [5, 16, 25] for details), of line sites under similar distance functions [6], and of sphere sites under the Euclidean metric [3]. Only very recently, the authors [18] have shown this to hold also in the case of arbitrary polyhedral sites under polyhedral distance functions.

In all the other, ‘open’ cases, cubic or near-cubic upper bounds for the complexity of 3-dimensional Voronoi diagrams are known. They are a consequence of the representation of such diagrams as lower envelopes of trivariate functions, each measuring the distance from a point in \mathbb{R}^3 to one of the sites; see [8] for this representation, and [24] for the bounds just stated. In contrast, only quadratic or near-quadratic lower bounds for the complexity of 3-dimensional diagrams are known [2, 6].

The case of the Euclidean metric appears to be harder than the case of polyhedral metrics (or distance functions), because the trivariate functions that measure distances are curved (except for the special case of point sites, where they can be transformed into linear functions), and the constraints that define the diagram are harder to analyze. The simplest open case of 3-dimensional Euclidean diagrams is that where the sites are lines. This specific problem is listed as Problem 3 in the list of open problems in computational geometry, recently published by Mitchell and O’Rourke [20]. A recent result that lends credence to the conjecture that the complexity of such diagrams is near-quadratic, is due to Agarwal and Sharir [1], who showed that the complexity of the union of n infinite congruent cylinders in 3-space is near-quadratic. The boundary of this union can be interpreted as a cross-section of the Euclidean Voronoi diagram of the axes of the cylinders, being the locus of all those points whose distance to the nearest axis has a fixed value (equal to the common radius of the cylinders). The complicated proof in [1], and the fact that the result applies merely to a single cross-section of the diagram, suggests that the problem involving the whole Euclidean Voronoi diagram of lines might be particularly hard to tackle.

Our contribution. In this paper we obtain the first result towards the described goal. We study the special case where the sites are lines that have a fixed number c of distinct orientations (and the metric is Euclidean). Even this special case is quite nontrivial to analyze. We show that the complexity of the diagram is $O(c^3 n^{2+\varepsilon})$,

for any $\varepsilon > 0$, where the constant of proportionality depends on ε . This implies, in particular, that when the number of distinct orientations in a collection of lines is constant (that is, $c = O(1)$), the complexity of its Euclidean Voronoi diagram is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This completely confirms the above-mentioned conjecture in this case.

Our result is structurally similar to that of Huttenlocher et al. [15], who have shown that the number of combinatorial changes in the Euclidean Voronoi diagram of c sets of rigidly moving points in the plane, each set of cardinality n , is close to $O(c^3 n^2)$.

Motivation The motivation underlying the study of Voronoi diagrams in computational geometry has always been algorithmic. They provide a natural data structure for handling a variety of applications, important both in theory and in practice, such as proximity (nearest neighbor) queries, high-clearance placements and motion planning problems, clustering and classification problems, and many more (see, among others, the survey by Aurenhammer and Klein [4] and the book by Okabe et al. [21] for a description of many of these applications).

There are several general techniques for computing Voronoi diagrams, such as randomized incremental construction or sweep-based methods, and many more ad-hoc approaches. However, a precursory stage to the design of any algorithm for computing Voronoi diagrams is obtaining sharp bounds on their complexity. This will serve as a lower bound for the efficiency of any such algorithm, and, quite often, can be used in the design of algorithms with roughly the same running time. Nevertheless, most of the algorithmic study of Voronoi diagrams has been confined to planar diagrams, for the good reason that we are still lacking sharp general bounds for the complexity of generalized 3-dimensional diagrams.

The results presented in this paper are an attempt to remedy this situation. The special case we treat is important because it provides us with one more problem instance where near-quadratic bounds can be established. We hope that the method developed here will find applications in the analysis of other types of 3-dimensional Voronoi diagrams (see the remark at the end of Section 3), and thereby lead us further towards the ultimate goal of establishing near-quadratic bounds for general 3-dimensional diagrams, following which near-quadratic algorithms for their construction will not be too difficult to design.

Moreover, the considered setting of lines with a fixed number of orientations is interesting in its own right. It is applicable, for example, to the problem of motion planning, or of finding largest free placements, of a ball amid a collection of ‘beams’ or ‘pipes’ in 3-space. It is a natural assumption that the beams have only a constant number of orientations. (Typical examples of this setting occur in architectural design.)

Organization. We first study, in Section 2, the special case where the lines have at most three distinct orientations. In this special case, we obtain the slightly improved bound $O(n\lambda_5(n))$, where $\lambda_5(n) = O(n \cdot \alpha(n)^{O(\alpha(n))})$ is the maximum length of Davenport-Schinzel sequences of order 5 on n symbols, and where $\alpha(n)$ is the extremely slowly growing inverse Ackermann function (see [24] for details). The case of four orientations is treated in Section 3, and the simple extension to more than four orientations is described in Section 4.

2 The Case of Two or Three Orientations

Let L be a set of n lines in 3-space, which have up to three distinct orientations. Thus L can be written as $R \cup B \cup G$, where all the lines in R (called ‘red’ lines) have the same orientation, and the same holds for the lines of B (‘blue’ lines) and those of G (‘green’ lines).

We adopt a limited general position assumption on L , as follows. First, we assume that each of the collections R , B , and G , is in general position, in the sense that its intersection with any fixed generic plane is a collection of points in general position (that is, it does not contain collinear triples or cocircular quadruples of points or other degenerate configurations). We also assume that the three vectors that are parallel to the orientations of the collections R , B , and G , do not lie in a common plane.

Before we proceed, we need to mention some basic properties of *bisectors* and *trisectors* of lines, which are, respectively, the loci of points equidistant from two and three lines. These geometric properties are reported here without proofs, which are given as an appendix below, in order to maintain the flow of exposition. The main conclusions from the analysis carried out in the appendix are as follows. A bisector of two lines is in general a hyperbolic paraboloid, which is a doubly ruled quadratic surface (it degenerates to a plane when the two lines are parallel). A trisector of three pairwise non-parallel lines is an algebraic curve of degree four, and, if non-singular, has exactly four components, all unbounded. If two of the three defining lines are parallel, the trisector becomes a planar conic section (of degree two, consisting of at most two unbounded components). If all the three lines are parallel, the trisector is a line parallel to them. The fact that no component of any trisector is bounded will be significant in our analysis. In the sequel, we will denote the bisector of two lines e, f by $H_{e,f}$, and the trisector of three lines e, f, g will be denoted by $\tau_{e,f,g}$.

Denote the Euclidean Voronoi diagram of L by $Vor(L)$. We begin by bounding the number of its vertices. Let v be such a vertex, incident to the cells of four lines $\ell_1, \ell_2, \ell_3, \ell_4$. At least two of them must be of the same color. Suppose first that three of them are of the same color, say, $\ell_1, \ell_2, \ell_3 \in R$. Project v and all the lines of R onto a plane π orthogonal to these lines. Then each line of R projects to a point, and v projects onto a vertex v^* of the planar Voronoi diagram of the projected points within π . The number of such vertex projections v^* is thus at most $2n - 4$. Moreover, the number of vertices v that can project onto the same point v^* is at most $2n$. This is

because the radius r of the ball centered at v and touching ℓ_1, ℓ_2, ℓ_3 is equal to the radius of the disk within π centered at v^* and touching the point projections of these three lines. As we slide a ball of radius r while maintaining contact with ℓ_1, ℓ_2, ℓ_3 , we reach at most $2n$ placements where it touches a fourth line. Each of these touching placements in which the ball is not crossed by any other line gives rise to a Voronoi vertex that projects onto v^* . This implies that the overall number of Voronoi vertices of the kind under consideration is at most $(2n - 4) \cdot 2n = O(n^2)$.

Suppose then that exactly two of the four lines are of the same color, say $\ell_1, \ell_2 \in R$, $\ell_3 \in B$, and $\ell_4 \in G$. If we project v and the lines of R onto the same plane π as above, we obtain that the projection of v lies on a Voronoi edge of the planar diagram of the point projections of the red lines. The number of such edges is $O(n)$.

Fix such an edge e , and consider the 2-dimensional slab Σ_e obtained by sweeping e in the direction of the red lines; by construction, $v \in \Sigma_e$. Moreover, Σ_e is the locus of all the centers of balls that touch ℓ_1 and ℓ_2 , and no other red line. Let H_e denote the plane containing Σ_e , and let ℓ_0 be the line of intersection between H_e and the plane π_0 spanned by ℓ_1 and ℓ_2 — this intersection is the midline of the 2-dimensional slab spanned by ℓ_1 and ℓ_2 . Denote the two halfspaces bounded by π_0 as π_0^+ and π_0^- . See Figure 1.




Figure 1: The bisector of ℓ_1 and ℓ_2 .

Fix a point $p \in \ell_0$, and consider the line λ_p that passes through p , lies in H_e , and is orthogonal to ℓ_0 . Parametrize λ_p by a real parameter y , where $y = 0$ at p , $y > 0$ within π_0^+ , and $y < 0$ within π_0^- . Move a point q along the entire λ_p , in the direction of increasing y . The ball centered at q and touching ℓ_1, ℓ_2 has the property that its intersection with π_0^+ keeps expanding during the motion (i.e., any point of π_0^+ that the moving ball meets will remain inside the ball as its center keeps moving in the above direction). Similarly, the portion of the moving ball within π_0^- keeps shrinking ‘into itself’.

Let each line $\ell \in B \cup G$ define two rays $\ell^+ = \ell \cap \pi_0^+$, $\ell^- = \ell \cap \pi_0^-$. With each ray ℓ^+ (resp., ℓ^-), associate a function ψ_{ℓ^+} (resp., ψ_{ℓ^-}) on ℓ_0 , where $\psi_{\ell^+}(p)$ (resp., $\psi_{\ell^-}(p)$), for $p \in \ell_0$, is the y -value of the center of the ball that touches ℓ_1, ℓ_2 , and ℓ^+ (resp., ℓ^-), where the center lies on λ_p . The functions $\psi_{\ell^+}, \psi_{\ell^-}$ are defined (and continuous) when ℓ does not intersect the disk centered at p , lying in π_0 , and touching ℓ_1 and ℓ_2 . Hence, the (common) domain of definition of ψ_{ℓ^+} and ψ_{ℓ^-} is either the full line ℓ_0 , if ℓ does not intersect the 2-dimensional slab spanned by ℓ_1 and ℓ_2 , or the union of two rays along ℓ_0 , otherwise.

Denote the collection of the functions ψ_{ℓ^+} (resp., ψ_{ℓ^-}), for $\ell \in B \cup G$, by Ψ^+ (resp., by Ψ^-). The preceding observations imply that any Voronoi vertex $v \in \Sigma_e$ under consideration (two of whose defining lines are in $B \cup G$) corresponds either to a vertex of the lower envelope of Ψ^+ , or to a vertex of the upper envelope of Ψ^- , or to an intersection point between the two envelopes.

It is easily seen that any pair of functions of the above kind intersect in at most four points. Indeed, any such intersection point w is equidistant from ℓ_1, ℓ_2 , and from two other lines $\ell_3, \ell_4 \in B \cup G$. That is, we have

$$d^2(w, \ell_1) \left(= d^2(w, \ell_2) \right) = d^2(w, \ell_3) = d^2(w, \ell_4).$$

The squared distance of a point w from a line that passes through a point a and has unit direction u , is

$$\|w - a\|^2 - ((w - a) \cdot u)^2,$$

which is a quadratic polynomial in the coordinates of w . Since w lies on the plane H_e , we obtain a system of two quadratic equations in two variables, which has at most four solutions (see also the proof of Lemma 3.1 below).

It is shown, e.g., in [24, Lemma 1.8] that the complexity of the upper or lower envelope of continuous functions, so that each function is defined on a ray or on the whole real line, and so that each pair of them intersect in at most four points, is $O(\lambda_5(n)) = O(n \cdot \alpha(n)^{O(\alpha(n))})$ [24], where $\lambda_5(n)$ is the maximum length of Davenport-Schinzel sequences of order 5 on n symbols, and where $\alpha(n)$ is the extremely slowly growing inverse Ackermann function. As observed above, we can split each partially defined function in $\Psi^+ \cup \Psi^-$ into two functions, each defined over a ray. We thus conclude that the number of Voronoi vertices in (the relative interior of) Σ_e is $O(\lambda_5(n))$. Multiplying this bound by the number $O(n)$ of edges e , and adding the preceding bound $O(n^2)$ on the number of vertices defined by three lines of the same color, we conclude that the number of vertices of the diagram $Vor(L)$ is $O(n\lambda_5(n))$.

We next bound the number of edges of $Vor(L)$. If an edge e is delimited by a Voronoi vertex v , we charge e to v . By the general position assumption, each v is charged at most 4 times, so the number of edges e of this kind is $O(n\lambda_5(n))$. Let e be a Voronoi edge that has no incident Voronoi vertex. As mentioned above, the analysis of trisectors implies that e is not bounded.

Fix two planes $\pi^\pm : z = \pm z_0$, such that each unbounded edge of $Vor(L)$ intersects at least one of them (assuming that the coordinate directions are generic, such planes

exist). It therefore suffices to bound the complexity of the cross sections of $Vor(L)$ with the planes π^\pm . Consider, say, the plane π^+ . The Voronoi cells in each of the monochromatic diagrams $Vor(R)$, $Vor(B)$, $Vor(G)$ are unbounded convex prisms, whose faces are all parallel to the orientation of the respective collection of lines, and the overall complexity of each diagram is $O(n)$. Hence, the intersection of π^+ with each of these monochromatic diagrams is a planar convex subdivision of complexity $O(n)$. The overlay of these cross sections is a planar convex subdivision of complexity $O(n^2)$. For each cell ξ of the overlay, there exists a fixed red line r , a fixed blue line b and a fixed green line g , which are the nearest red, blue, and green lines to any point in ξ , respectively. It follows that the complexity of the overall diagram $Vor(L)$ within ξ is bounded by a constant, which implies that the complexity of the diagram within π^+ (and, symmetrically, within π^-) is $O(n^2)$.

This implies that the number of unbounded edges of $Vor(L)$ is $O(n^2)$. It is easily seen that the number of 2-faces of the diagram is proportional to the number of vertices plus the number of edges plus $O(n^2)$. Finally, the number of 3-cells is only n : Each line has a connected, star-shaped Voronoi cell [19]. Hence, we obtain the main result of this section:

Theorem 2.1. *The complexity of the Voronoi diagram of a set of n lines with at most three distinct orientations is $O(n\lambda_5(n)) = O(n^2 \cdot \alpha(n)^{O(\alpha(n))})$.*

3 The Case of Four Orientations

We now assume that the given set L of lines is the union of four subsets, each consisting of lines at a fixed direction. We denote these subsets by R (consisting of ‘red’ lines), B (consisting of ‘blue’ lines), G (consisting of ‘green’ lines), and Y (consisting of ‘yellow’ lines). The proof of the following elementary geometric fact is provided for completeness.

Lemma 3.1. *The maximum number of balls tangent to four given lines in 3-space, assuming general position, is 8.*

Proof. As already noted, the distance $d(\mathbf{x}, \ell)$ between a point $\mathbf{x} \in \mathbb{R}^3$ and a line ℓ , passing through a point a and having unit direction u , satisfies

$$d^2(\mathbf{x}, \ell) = \|\mathbf{x} - a\|^2 - ((\mathbf{x} - a) \cdot u)^2,$$

which is a quadratic function of \mathbf{x} . Given four lines $\ell_1, \ell_2, \ell_3, \ell_4$ in general position, the center \mathbf{x} of a ball that is tangent to all four lines has to satisfy the equations

$$d^2(\mathbf{x}, \ell_1) = d^2(\mathbf{x}, \ell_2) = d^2(\mathbf{x}, \ell_3) = d^2(\mathbf{x}, \ell_4).$$

These are three quadratic equations, so, by Bezout’s Theorem [14], the number of solutions is at most $2^3 = 8$.

The number 8 can be attained: We first give a construction where the lines are not in general position. Take ℓ_1, ℓ_2, ℓ_3 to be any three nonconcurrent lines in the xy -plane. They determine four disks D_1, D_2, D_3, D_4 in that plane that are tangent to all three of them, as shown in Figure 2. Take ℓ_4 to be any line perpendicular to the xy -plane, meeting the plane at a point not lying in any of these disks. Fix a disk D_i , and let λ_i be the z -vertical line passing through the center of D_i ; this is the locus of all centers of balls that touch ℓ_1, ℓ_2, ℓ_3 and meet the xy -plane at D_i . It is easily seen that there are exactly two points on λ_i , symmetric to each other with respect to the xy -plane, that are centers of balls that also touch ℓ_4 . For any specific disc D_i , this yields two distinct balls that touch all four lines, giving us 8 such balls overall. By slightly perturbing the lines, we can obtain a construction for lines in general position. This completes the proof of the lemma. \square




Figure 2: Four lines having eight Voronoi vertices.

Let $\ell_1, \ell_2, \ell_3, \ell_4$ be four given lines of different colors. Let $s \leq 8$ denote the number of balls tangent to all four of them, and let c_1, \dots, c_s denote the centers of these balls, sorted in increasing order of their x -coordinate. (The coordinate frame is assumed to be generic, so that no two c_i 's have the same x -coordinates.) Define the *index* $\text{ind}(c_i)$ of c_i to be $\min\{i - 1, s - i\}$, so we have $0 \leq \text{ind}(c_i) \leq 3$ for each i .

With each line $\ell \in L = R \cup B \cup G \cup Y$ we associate the squared distance function $f_\ell : \mathbb{R}^3 \mapsto \mathbb{R}$, given by $f_\ell(\mathbf{x}) = d^2(\mathbf{x}, \ell)$. Let $\mathcal{E}_{\mathcal{F}}$ denote the lower envelope of the set $\mathcal{F} = \mathcal{F}(L) = \{f_\ell \mid \ell \in L\}$. Clearly, the *minimization diagram* of $\mathcal{E}_{\mathcal{F}}$, namely, the projection of (the graph of) $\mathcal{E}_{\mathcal{F}}$ onto the xyz -space, is the Voronoi diagram $\text{Vor}(L)$ (see also [8]).

For each point $\mathbf{q} = (q_1, q_2, q_3, q_4) \in \mathbb{R}^4$, define its *R-level* (resp., *B-level*, *G-level*, *Y-level*) to be the number of lines $\ell \in R$ (resp., $\ell \in B$, $\ell \in G$, $\ell \in Y$) whose corresponding function graphs pass below \mathbf{q} ; that is, $q_4 > f_\ell(q_1, q_2, q_3)$. The *combined level* of \mathbf{q} is the sum of its red, blue, green and yellow levels. We denote the graph of each $f_\ell \in \mathcal{F}$ by \tilde{f}_ℓ . Denote by \tilde{R} the collection of all graphs \tilde{f}_ℓ , for $\ell \in R$, and define \tilde{B} , \tilde{G} , \tilde{Y} , and \tilde{L} , analogously. Let $\mathcal{A}(\tilde{L})$ denote the arrangement in \mathbb{R}^4 of the graphs

\tilde{f}_ℓ of the functions in \tilde{L} . Clearly, for a vertex \mathbf{q} of $\mathcal{A}(\tilde{L})$, \mathbf{q} is a vertex of $\mathcal{E}_{\mathcal{F}}$ if and only if the combined level of \mathbf{q} is 0.

Let $V_0^{(j)}(L)$ (resp., $V_{\leq k}^{(j)}(L)$) denote the number of ‘4-colored’ vertices \mathbf{q} of $\mathcal{A}(\tilde{L})$ (i.e., vertices incident to a red graph, a blue graph, a green graph and a yellow graph) of index $\leq j$, whose combined level is 0 (resp., at most k). Put $V_0(L) = V_0^{(3)}(L)$ and $V_{\leq k}(L) = V_{\leq k}^{(3)}(L)$. We also put $V_0^{(j)}(n) = \max_L V_0^{(j)}(L)$, where the maximum is taken over all families L of n lines, each having one of the four given orientations; $V_{\leq k}^{(j)}(n)$ is defined analogously. Using the Clarkson-Shor bound on levels [7], we have

$$V_{\leq k}^{(j)}(n) = O\left(k^4 V_0^{(j)}\left(\frac{n}{k}\right)\right).$$

As mentioned in Section 2 and proven in the appendix, every connected component of any trisector is unbounded. However, in the proof below we will *not* make use of this property at all. This will be significant when we extend the analysis to more general setups — see a discussion at the end of this section.

3.1 Irregular Vertices

Let v be a 4-colored vertex of the diagram, interpreted as a vertex of the lower envelope $\mathcal{E}_{\mathcal{F}}$, incident to four graphs $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$, for some $r \in R, b \in B, g \in G$ and $y \in Y$. The vertex v is incident to four edges of the envelope, which we denote mnemonically as rbg, rby, rgy , and bgy , where $rbg \subseteq \tau_{r,b,g}$ denotes the edge lying on the graphs $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g$, and similarly for the three other edges. As noted in [23], at least one of these edges emanates from v in the positive x -direction, and at least one edge emanates in the negative x -direction. We call v a *regular* vertex if exactly two of these edges emanate from v in the positive x -direction and exactly two in the negative x -direction. Otherwise, we call v *irregular*.

Lemma 3.2. *There are only $O(n\lambda_5(n))$ irregular vertices.*

Proof. Let v be an irregular vertex. If v is not 4-colored then the claim follows from Theorem 2.1, so assume that v is 4-colored, and use the above notations to denote the surfaces and edges incident to v . Suppose, without loss of generality, that three of the incident edges emanate from v to the left, and assume that they are rbg, rby , and rgy . In this case (assuming general position), v is a locally x -maximal vertex of the Voronoi cell $V(r)$ of r . Clearly, each line has a single connected Voronoi cell. In fact, each cell, being star-shaped with respect to its defining line, is also simply connected; see, e.g., [19].

As shown, e.g., in [12, Lemma 2.4], the number of locally x -extremal points of a simply connected 3-dimensional region K is proportional to 1 plus the number of *critical points* of ∂K (relative to the x -direction). These are points w for which the cross section of the interior of K with the yz -parallel plane through w is disconnected near w , but becomes connected (near w) when the plane slightly translates in some

direction. Hence, the number of irregular vertices of $Vor(L)$ is proportional to the number of critical points of cell boundaries plus $O(n)$.

Assuming general position, each critical point w of $\partial V(r)$ is incident to only 3 surfaces; it is typically a locally x -extremal point of a Voronoi edge of $V(r)$. Suppose, without loss of generality, that w is incident to $\tilde{f}_r, \tilde{f}_{b_1}, \tilde{f}_{g_1}$, for some $b_1 \in B, g_1 \in G$. Then w is a locally x -extreme point of (the relative interior of) a Voronoi edge (a portion of τ_{r,b_1,g_1}) of the 3-colored Voronoi diagram $Vor(R \cup B \cup G)$. By Theorem 2.1, the overall number of such features is $O(n\lambda_5(n))$, and this completes the proof of the lemma. \square

3.2 The Counting Scheme

In light of Lemma 3.2, this section is devoted to bounding the number of regular vertices of $Vor(L)$. This number is estimated using a variation of the ‘counting scheme’ technique, as introduced by Halperin and Sharir [11, 23] (see also [24]).

Let v be a 4-colored regular vertex, incident to $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$, using the notation introduced above. Let $0 \leq j \leq 3$ be the index of v . Without loss of generality, assume that there are exactly j vertices incident to $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$ to the *right* (that is, in the x -increasing direction) of v . By definition, v is incident to two edges of $\mathcal{E}_{\mathcal{F}}$ that emanate from it to the right, and to two edges that emanate from it to the left. Without loss of generality, assume that the edges emanating to the right are rbg and rby , and the edges emanating to the left are rgy and bgy .

Consider the 2-dimensional bisector $H_{g,y}$. Denote by R_{gy} the set of trisectors $\tau_{g,y,r'}$ drawn as curves along $H_{g,y}$, for red lines $r' \in R$. Define in an analogous manner the sets B_{gy}, G_{gy} , and Y_{gy} (where the latter two sets exclude the ill-defined trisectors induced by g and y themselves). Let \mathcal{A}_{gy} denote the 2-dimensional arrangement of the collection $R_{gy} \cup B_{gy} \cup G_{gy} \cup Y_{gy}$ of curves within $H_{g,y}$. It follows that there exists a face of \mathcal{A}_{gy} that is also a 2-face of $\mathcal{E}_{\mathcal{F}}$ on $H_{g,y}$, such that v is a locally x -maximal vertex of that face.

Let $\gamma_r \in R_{gy}$ (resp., $\gamma_b \in B_{gy}$) denote the trisector $\tau_{r,g,y}$ (resp., $\tau_{b,g,y}$), regarded as a curve within $H_{g,y}$. If we follow γ_r from v to the right, we lie, locally near v , above $\mathcal{E}_{\mathcal{F}}$ (actually, above \tilde{f}_b), and similarly for γ_b (which lies locally above \tilde{f}_r). See Figure 3.




Figure 3: The vertex v on $H_{g,y}$ —a view from the bottom (in \mathbb{R}^4).

3.2.1 Initial counting stages and vertices of index 0 and 1.

Lemma 3.3. $V_0^{(0)}(n)$ and $V_0^{(1)}(n)$ are bounded by $O(n\lambda_5(n))$.

Proof. Trace the curve γ_r from v to the right, and stop as soon as we reach one of the following *critical events* along γ_r :

- (a) We reach another intersection of the four graphs $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$.
- (b) We reach a 3-colored vertex.
- (c) We reach $x = +\infty$.
- (d) We reach a locally x -extremal point of the curve γ_r .

We refer to events of types (b)–(d) as *terminal events*.

Perform a similar tracing along γ_b . Suppose that at least one of the tracings, say, along γ_r , reaches a terminal event. The first such event is either a vertex of the 3-colored Voronoi diagram of $R \cup G \cup Y$, or it can be charged to an edge of this diagram. By Theorem 2.1, the number of such events is thus $O(n\lambda_5(n))$, and each such event is uniquely counted by some vertex v (this follows since between v and the terminal event we are always above $\mathcal{E}_{\mathcal{F}}$). Hence, the number of vertices v that fall in this case is $O(n\lambda_5(n))$. In particular, this bounds the number of vertices of index 0.

We may thus assume that the tracing of γ_r ends at a vertex u , and the tracing of γ_b ends at a vertex w , so that both u and w are incident to $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$ (see Figure 4). Moreover, the portion $\delta_r^{(1)}$ of γ_r between v and u , and the portion $\delta_b^{(1)}$ of γ_b between v and w , are both x -monotone, and neither of them contains a 3-colored vertex or another terminal event. In particular, u and w lie to the right of v , the red, green and yellow levels of u are all 0, and the blue, green and yellow levels of w are all 0.




Figure 4: Tracing from v to the right.

If $u = w$ then this is a vertex of the diagram (because all its colored levels are 0), of index at most $j - 1$ (because it lies to the right of v). The number of vertices v in this subcase is thus at most $V_0^{(j-1)}(n)$. In particular, this is easily seen to imply that the number of vertices of index 1 is $O(n\lambda_5(n))$. We have thus showed the following:

$$V_0^{(0)}(n) = O(n\lambda_5(n)), \quad V_0^{(1)}(n) = O(n\lambda_5(n)),$$

which wraps up the proof of the lemma. \square

3.2.2 Subsequent counting stages and vertices of index 2.

In what follows, we assume that $j = 2$ or 3 . In light of the arguments made in the proof of Lemma 3.3, we may assume that $u \neq w$. Fix some threshold parameter k , to be determined later.

Lemma 3.4. $V_0^{(2)}(n)$ is bounded by $O(k^3 n \lambda_5(n) + k^2 V_0(\frac{n}{5k}))$.

Proof. Suppose that the blue (and thus the combined) level of u is at most $4k$. In this case, we charge v to u . The charging is unique, implying that the number of vertices v in this case is at most

$$V_{\leq 4k}^{(j-1)}(n) = O\left(k^4 V_0^{(j-1)}\left(\frac{n}{4k}\right)\right),$$

where, as already mentioned, we use the Clarkson-Shor bound on levels [7]. A similar charging is applied if the level of w is at most $4k$. Hence, in what follows, we may assume that $u \neq w$ and that both lie at combined level $> 4k$.

Let W denote the portion of $H_{g,y}$ consisting of all points that lie above the graphs of both f_r and f_b , and let W_0 be the connected component of W whose boundary contains v . The region W_0 is bounded, locally near v and to its right, by the two arcs $\delta_r^{(1)}$ and $\delta_b^{(1)}$, and v is a locally leftmost (x -minimal) vertex of W_0 . Let $\delta_b^{(2)}$ (resp., $\delta_r^{(2)}$) denote the other edge of ∂W_0 incident to u (resp., to w). Both $\delta_r^{(1)}$ and $\delta_r^{(2)}$ are contained in the trisector $\tau_{r,g,y}$, although they do not have to lie on the same component of that curve. Similarly, $\delta_b^{(1)}$ and $\delta_b^{(2)}$ are contained in $\tau_{b,g,y}$. Without loss of generality, we assume that $\delta_r^{(1)}$ lies clockwise to $\delta_b^{(1)}$ (when viewed from above); see Figure 5 for an illustration of several possible shapes of W_0 .

Let ζ be a vertex of \mathcal{A}_{gy} along $\delta_r^{(1)}$, incident to the graph of some other blue function f_β . (Recall that, by assumption, all vertices along $\delta_r^{(1)}$ are 4-colored). Consider the trisector $\tau_{\beta,g,y}$ as a curve γ_β within $H_{g,y}$, and let δ_β denote the connected component of $\gamma_\beta \cap W_0$ incident to ζ . We say that δ_β is a *deep* arc, if it contains at least k vertices of \mathcal{A}_{gy} . If δ_β is not deep, and it contains a terminal event (namely, it contains a 3-colored vertex, or contains a locally x -extremal point, or reaches $x = \pm\infty$), we call it a *terminal* arc. Otherwise, we call it *shallow*. See Figure 6 for a special case of a shallow arc. Similar notations apply to red arcs that emanate from vertices of \mathcal{A}_{gy} along $\delta_b^{(1)}$.

Consider the first $4k$ vertices along $\delta_r^{(1)}$ (by assumption, $\delta_r^{(1)}$ must contain at least these many vertices). If at least $2k$ of the corresponding arcs δ_β are deep, then collecting the first k vertices along each of these arcs yields a set of at least $2k^2$ vertices of \mathcal{A}_{gy} within W_0 , all lying at combined level at most $5k$. We claim that each of them is charged by vertices like v at most a constant number of times. Indeed, let η be such a vertex, lying on a deep blue arc δ_β . Note that the starting point ζ of




Figure 5: Several possible structures of the region W_0 . In all cases, at most 3 vertices of W_0 lie to the right of v .




Figure 6: A shallow arc that lands back on $\delta_r^{(1)}$.

δ_β is at red level 0, but all points in the relative interior of δ_β have strictly positive red levels. Hence, we can trace δ_β back from η (there are two possible directions for this tracing) until we reach the first point ζ at red level 0. The point ζ must lie on $\delta_r^{(1)}$, and we can trace $\delta_r^{(1)}$ from ζ backwards (to the left) until we reach v —the first vertex at combined level 0. Hence, using [7], as above, the number of vertices v in this subcase is at most

$$O\left(\frac{1}{k^2}V_{\leq 5k}(n)\right) = O\left(k^2V_0\left(\frac{n}{5k}\right)\right).$$

The same bound applies to the number of vertices v for which at least $2k$ of the first $4k$ vertices along $\delta_b^{(1)}$ are sources of deep red arcs.

Hence, we may assume that among the first $4k$ vertices along $\delta_r^{(1)}$, at least $2k$

are sources of shallow or terminal arcs, and similarly for $\delta_b^{(1)}$. If any of these arcs is terminal, we charge v to the corresponding terminal event along the arc. We note that such an event η lies at combined level at most $5k$. Hence, η is or can be charged to an ($\leq 5k$)-level feature of one of the 4-dimensional 3-colored arrangements $\mathcal{A}(\tilde{B} \cup \tilde{G} \cup \tilde{Y})$, $\mathcal{A}(\tilde{R} \cup \tilde{G} \cup \tilde{Y})$. Moreover, arguing as in the preceding paragraph, η is charged by vertices like v at most twice. By Theorem 2.1 and [7], the number of such events η , and thus also the number of vertices v that fall into this subcase, is at most

$$O\left(k^4 \cdot \frac{n}{5k} \lambda_5\left(\frac{n}{5k}\right)\right) = O(k^2 n \lambda_5(n)).$$

Hence, we may assume that at least $2k$ of the first $4k$ vertices along $\delta_r^{(1)}$ are sources of shallow arcs, and none of these vertices are sources of terminal arcs. Moreover, the same property holds for $\delta_b^{(1)}$.

Suppose that one of these shallow arcs, δ_β , emanating from $\delta_r^{(1)}$, terminates also on $\delta_r^{(1)}$, as in Figure 6. By definition, δ_β does not encounter any blue graph $\tilde{f}_{\beta'}$ (for then δ_β would contain a 3-colored vertex and thus would be terminal). Hence, the blue level of the terminal endpoint ζ' of δ_β is equal to the blue level of the starting point ζ , and all other levels of both endpoints are 0. In this case we skip the portion of $\delta_r^{(1)}$ between ζ and ζ' . More precisely, we modify the tracing procedure used so far as follows: Trace $\delta_r^{(1)}$ to the right, starting from v , and attempt to collect either $2k$ deep arcs, or a terminal arc, or $2k$ shallow arcs that do not terminate on $\delta_r^{(1)}$. If during this tracing we reach a shallow arc δ_β that does terminate on $\delta_r^{(1)}$, we take a ‘shortcut’ along δ_β , and continue the tracing of $\delta_r^{(1)}$ from the other endpoint of δ_β . It is clear that this modified process must terminate successfully, or else we would reach the endpoint u of $\delta_r^{(1)}$, which then would lie at level $\leq 4k$, contrary to assumption. From now on, we apply a similarly modified tracing procedure to $\delta_b^{(1)}$ as well.

We thus reach the following situation. We have collected at least $2k$ shallow blue arcs that emanate from $\delta_r^{(1)}$ and terminate on other red edges of ∂W_0 , and at least $2k$ shallow red arcs that emanate from $\delta_b^{(1)}$ and terminate on other blue edges of ∂W_0 . The combined level of any point on any of these arcs is at most $5k$.

Suppose that one of the shallow blue arcs δ_β that emanates from $\delta_r^{(1)}$ terminates on a (red) edge $\delta_r^{(3)}$ of ∂W_0 that does not intersect $\tau_{b,g,y}$ at all. That is, $\delta_r^{(3)}$ is a full (bounded or unbounded) component of the trisector $\tau_{r,g,y}$, which lies fully above the graph of f_b , as in Figure 7. Let η be the ‘landing point’ of δ_β on $\delta_r^{(3)}$. The combined level of η is at most $5k$. Trace $\delta_r^{(3)}$ from η to the right (i.e., in the positive x -direction) until we reach a terminal event η' , to which we charge v . (Such an η' always exists: even if we do not encounter any finite event, we will reach $x = +\infty$, which is a terminal event, by definition.) Note that η' is or can be charged to a feature of the 4-dimensional arrangement $\mathcal{A}(\tilde{R} \cup \tilde{G} \cup \tilde{Y})$, whose combined level (in this 3-colored arrangement) is at most $5k$. Arguing as above, the number of such events is $O(k^2 n \lambda_5(n))$. Here we cannot claim that η' is uniquely charged by v , but we can still bound the number of times η' is charged, as follows.

Trace $\delta_r^{(3)}$ back (in the negative x -direction) from η' (there may be two choices




Figure 7: Charging v when a shallow arc lands on an edge of W_0 that does not meet other such edges.

for r and for $\delta_r^{(3)}$ given a specific η' , since η' may be a 3-colored vertex), until the first time we reach a point whose combined level (including the blue level) is at most $5k$. This backwards tracing has to succeed: it will reach η or stop earlier. Then any charging vertex v must be a vertex incident to $\tilde{f}_r, \tilde{f}_g, \tilde{f}_y$, and to some \tilde{f}_{b_1} , from the at most $5k$ blue surfaces b_1 that lie below the stopping point. In other words, η' can be charged at most $O(k)$ times, implying that the number of vertices v that fall into this subcase is

$$O(k^3 n \lambda_5(n)). \quad (1)$$

A symmetric analysis applies if a shallow red arc lands on a blue edge of ∂W_0 which is a full component of $\tau_{b,g,y}$. Moreover, the analysis just given also holds if $\delta_r^{(3)}$ meets \tilde{f}_b in only one of the two directions from η , and extends to infinity in the other direction. It also holds if, in at least one of the two directions, we meet a terminal event before meeting \tilde{f}_b . And it also holds in the symmetric extended cases where the roles of the red and blue colors are interchanged.

The above analysis implies, in particular, that in the sequel, we may assume that none of the first $2k$ shallow arcs that emanate from $\delta_r^{(1)}$ and $\delta_b^{(1)}$ terminate on a bounded component of ∂W_0 that does not meet other components of ∂W_0 . Note also that the analysis holds if $\delta_r^{(3)}$ is a bounded component of $\tau_{r,g,y}$ that lies fully to the right of v . Indeed, even if such a component does meet \tilde{f}_b , it must meet it at two points, both different from u, w and lying to the right of v , which is impossible.

Suppose now that one of the collected blue shallow arcs δ_β terminates on $\delta_r^{(2)}$, as in Figure 8. Each of the $\geq 2k$ red shallow arcs that we have collected along $\delta_b^{(1)}$ must cross δ_β . Indeed, none of these arcs terminate on $\delta_b^{(1)}$, by construction; they cannot terminate on $\delta_r^{(1)}$ or on $\delta_r^{(2)}$, for that would have made them terminal; and, as argued above, they also do not terminate on an isolated bounded component of ∂W_0 . This, however, contradicts the shallowness of δ_β , since it cannot contain more than k crossings with other arcs. We have thus showed that none of the collected blue shallow arcs terminate on $\delta_r^{(2)}$. Symmetrically, it can be shown that none of the

collected red shallow arcs terminate on $\delta_b^{(2)}$.




Figure 8: A shallow blue arc cannot ‘intercept’, by terminating on $\delta_r^{(2)}$, the shallow red arcs emanating from $\delta_b^{(1)}$.

The bounds accumulated so far account for all the vertices v with index at most 2. Specifically, we have

$$V_0^{(2)}(n) = O\left(k^3 n \lambda_5(n) + k^2 V_0\left(\frac{n}{5k}\right)\right),$$

thereby proving the lemma. □

3.2.3 Final counting stages and vertices of index 3.

Lemma 3.5. $V_0^{(3)}(n)$ is bounded by

$$O\left((k^2 \ell^2 + \ell^3) n \lambda_5(n) + V_0^{(2)}(n) + k^4 V_0^{(2)}\left(\frac{n}{4k}\right) + k^2 V_0^{(3)}\left(\frac{n}{5k}\right) + \ell^4 V_0^{(2)}\left(\frac{n}{5\ell}\right) + k^2 \ell^2 V_0^{(3)}\left(\frac{n}{6\ell}\right)\right).$$

Proof. From now on, we deal with vertices v of index 3. They are treated by considering a number of possible structures of the region W_0 as well as possible behavior patterns of arcs inside W_0 , and bounding the maximal number of vertices v in each case. This will often be performed by charging v to certain features in W_0 .

We already have sufficient machinery to dispose of vertices v for which $\delta_r^{(2)}$ and $\delta_b^{(2)}$ meet at a common endpoint, as in Figure 9. The preceding arguments allow us to assume that there are no shallow arcs that connect $\delta_r^{(1)}$ to $\delta_r^{(2)}$, or $\delta_b^{(1)}$ to $\delta_b^{(2)}$, and that there are no shallow arcs that land on any bounded component of W_0 within the quadrangle formed by $\delta_r^{(1)}$, $\delta_b^{(1)}$, $\delta_r^{(2)}$, and $\delta_b^{(2)}$. This means that in this case, unless u and w have level $O(k)$, we can either collect a terminal arc or at least $2k$ deep arcs, when sliding from v as above. In other words, we can charge v (almost uniquely) either to $\Theta(k^2)$ low-level vertices (at level $O(k)$) within W_0 , or to a low-level terminal event within W_0 , or to some other vertex of W_0 (that is, to u or to w), which lies at

level at most $4k$ and has a smaller index. Arguing as above, the number of vertices v that fall into this case is

$$O\left(k^3 n \lambda_5(n) + k^2 V_0\left(\frac{n}{5k}\right) + k^4 V_0^{(2)}\left(\frac{n}{4k}\right)\right).$$

We may thus assume that $\delta_r^{(2)}$ and $\delta_b^{(2)}$ do not meet.




Figure 9: The case where W_0 is a quadrangle.

Suppose, without loss of generality, that the vertex u lies to the left of w . Then any shallow blue arc δ_β that emanates from $\delta_r^{(1)}$ must terminate at a point that lies to the right of v (regardless of whether it extends to the right or to the left); see Figure 10. This is due to the fact that these arcs are x -monotone. Let η be the terminal point of δ_β , and let $\delta_r^{(3)} \neq \delta_r^{(1)}, \delta_r^{(2)}$ denote the red edge of ∂W_0 that contains η . (The preceding analysis implies that we may assume that all shallow blue arcs that we have collected do land on a new red edge of ∂W_0 .) Trace $\delta_r^{(3)}$ from η in the increasing x -direction. (Note the two different situations that can arise, where we can turn from δ_β to the traced portion of $\delta_r^{(3)}$ either to the left or to the right.¹) By the analysis just given, we may assume that this portion of $\delta_r^{(3)}$ terminates at a vertex t of W_0 , incident to $\tilde{f}_r, \tilde{f}_b, \tilde{f}_g, \tilde{f}_y$, which lies to the right of v and is different from u, w . That is, t is the ‘missing’ third sibling vertex of v that lies to the right of v . Moreover, the portion of $\delta_r^{(3)}$ between η and t contains no terminal event.

Suppose first that w lies to the left of t ; see Figure 11. There must exist a red shallow arc δ_ρ that emanates from $\delta_b^{(1)}$ and does not cross δ_β (and it also cannot cross $\delta_r^{(1)}, \delta_r^{(2)}$, and $\delta_r^{(3)}$). Since δ_ρ is x -monotone, it must terminate at a point η' to the right of v , regardless of whether it extends to the right or to the left: the concatenation of $\delta_r^{(1)}, \delta_\beta$, and $\delta_r^{(3)}$ up to t , does not allow δ_ρ to reach points left of v , because t lies to the right of w ; see Figure 11. The point η' lies on some blue edge $\delta_b^{(3)} \neq \delta_b^{(1)}, \delta_b^{(2)}$. Tracing $\delta_b^{(3)}$ from η' in the positive x -direction, we may assume that it terminates at a vertex of W_0 (the case of a terminal event can be charged as above), which is necessarily t itself. Moreover, the portion of $\delta_b^{(3)}$ between η' and t contains no terminal event. We now note that the red and blue levels of t are both at most k , since the red level of η and the blue level of η' are at most k , and since there are no terminal events on $\delta_r^{(3)}$

¹Recall that in the analysis of W_0 we refer to the view of this region from above (in the vertical direction of \mathbb{R}^4).




Figure 10: If u lies to the left of w , a shallow arc emanating from $\delta_r^{(1)}$ must terminate to the right of v . In (a) we make a left turn from δ_β to $\delta_r^{(3)}$ at η , and in (b) we make a right turn. In both cases $\delta_r^{(3)}$ has to contain a vertex t of W_0 in the direction of our tracing.

between η and t , and on $\delta_b^{(3)}$ between η' and t . Thus, the combined level of t is $O(k)$. Since t is of index at most 1 (it lies to the right of w), and is uniquely charged by v , the number of vertices v in this subcase is $O(k^2 n \lambda_5(n))$.




Figure 11: The case where w lies to the left of t .

Suppose then that w lies to the right of t . If any shallow red arc that emanates from $\delta_b^{(1)}$ and does not cross δ_β terminates to the right of v , we proceed as in the case, just treated, where t lies to the right of w . The only way in which this does not occur is when all these shallow red arcs emanate from $\delta_b^{(1)}$ in the negative x -direction, starting to the right of t and ‘bypassing’ the concatenation of $\delta_r^{(1)}$, δ_β , and $\delta_r^{(3)}$ up to t . See Figure 12.

To handle this case, choose another threshold parameter $\ell \gg k$, to be determined later. If t lies at level at most $5k + 4\ell$, we charge v to t . We note, as above, that the charging is unique, and use the fact that the index of t is at most 2, to conclude that the number of vertices v in this subcase is at most $V_{\leq 5k+4\ell}^{(2)}(n)$. Our choice of ℓ will




Figure 12: The case where w lies to the right of t and at a high level. The figure depicts the subcase where there are many deep blue arcs emanating from $\bar{\delta}_r^{(3)}$.

ensure that $5k + 4\ell \leq 5\ell$, so, using [7], the number of vertices v under consideration is

$$O\left(\ell^4 V_0^{(2)}\left(\frac{n}{5\ell}\right)\right).$$

Assume then that the level of t is $> 5k + 4\ell$. Then the portion $\bar{\delta}_r^{(3)}$ of $\delta_r^{(3)}$ between η and t must contain at least 4ℓ (4-colored) vertices. We now apply a collection process for blue arcs that emanate from $\bar{\delta}_r^{(3)}$. The process is very similar to that applied to $\delta_r^{(1)}$ (and to $\delta_b^{(1)}$), except that we redefine the notions of being deep, terminal, or shallow, in terms of the parameter ℓ rather than k . To distinguish between the old and new notions, we say that an arc is ℓ -deep (resp., ℓ -shallow) if it contains at least ℓ (resp., fewer than ℓ) vertices (and so that none of the first ℓ vertices is terminal). If one of the first ℓ vertices lying on an arc is terminal, the arc is said to be ℓ -terminal. The old notions are from now on designated, in complete analogy, as k -deep, k -shallow, and k -terminal.

The collection process on $\bar{\delta}_r^{(3)}$ is therefore as follows. Starting from η , we proceed along $\bar{\delta}_r^{(3)}$, taking shortcuts along ℓ -shallow arcs that land back on $\bar{\delta}_r^{(3)}$, and collect either 2ℓ ℓ -deep blue arcs, or an ℓ -terminal blue arc, or 2ℓ ℓ -shallow blue arcs that do not terminate on $\bar{\delta}_r^{(3)}$. The starting point of any collected arc is at blue level at most $4k + 4\ell \leq 5\ell$, and at red level at most k .

A significant technical difference between the two collection processes is that in the new process, we do not have the unique charging property that was utilized in the preceding analysis. Nevertheless, we do have a weaker property that we detail next.

Suppose that we have collected 2ℓ ℓ -deep blue arcs, as just described. See Figure 12. We thus obtain $\Theta(\ell^2)$ vertices along these arcs, all contained in W_0 and lying at combined level $4k + 5\ell \leq 6\ell$. We claim that each such vertex q is collected in this

fashion by at most $O(k^2)$ vertices v .

Consider such a vertex q , and attempt to trace back from q to determine the charging vertex v , as follows. Proceed from q along the ℓ -deep blue arc $\delta_{\beta'}$ that contains q , until the first time we reach a vertex q' at red level $\leq k$ (this will happen either when we reach $\bar{\delta}_r^{(3)}$ or earlier). The red surface incident to any charging vertex v must be one of the $\leq k$ red graphs that lie below q' (clearly, \tilde{f}_r is one of these graphs). Pick any of these graphs, $\tilde{f}_{r'}$, and continue to trace $\delta_{\beta'}$ backwards until the first time it actually intersects $\tilde{f}_{r'}$. If the stopping point is at red level $> k$, then r' is a wrong guess. We thus keep picking candidate graphs in this fashion, until for one of them, $\tilde{f}_{r''}$, the backward-tracing of $\delta_{\beta'}$ reaches $\tilde{f}_{r''}$ at red level $\leq k$. Once this situation is attained, we trace the red curve $\gamma_{r''}$ that we have hit, in the negative x -direction, until the first time we reach a point ν whose blue level is at most $4k$ (as above, if no such point exists, then r'' is a wrong guess, and we keep trying with different candidates $\tilde{f}_{r''''}$). The blue arc incident to a charging vertex v that is incident to $\tilde{f}_{r''}$ must then correspond to one of the $\leq 4k$ blue graphs lying below ν (\tilde{f}_b is clearly one of them when $r'' = r$). Now note that knowing which red and blue arcs are incident to v determines v uniquely. We have thus shown that there are only $O(k^2)$ possible vertices v that can charge q . Hence, using [7], the number of vertices v in this subcase is

$$O(k^2) \cdot O\left(\frac{1}{\ell^2} V_{\leq 6\ell}(n)\right) = O\left(k^2 \ell^2 V_0\left(\frac{n}{6\ell}\right)\right).$$

Similarly, if we collect an ℓ -terminal blue arc, the terminal event along it is charged by at most $O(k^2)$ vertices v , and there are $O(\ell^2 n \lambda_5(n))$ such events. The number of vertices v in this subcase is thus $O(k^2 \ell^2 n \lambda_5(n))$.

We are left to treat the case in which we have collected 2ℓ ℓ -shallow blue arcs. Note that their starting points on $\bar{\delta}_r^{(3)}$ are at combined level at most $5k + 4\ell \leq 5\ell$. Trace any such arc $\delta_{\beta'}$ to its end-point η' , which lies on some red edge of ∂W_0 , and at combined level $\leq 6\ell$. Several cases can arise, as depicted in Figure 13:

(a) $\eta' \in \delta_r^{(1)}$ and we make a *right* turn from $\delta_{\beta'}$ to $\delta_r^{(3)}$ at η : See Figure 13(a). In this case we trace $\delta_r^{(3)}$ from η to the left (in the negative x -direction). Since η lies to the left of w , it is easily seen that this tracing of $\delta_r^{(3)}$ must reach a local x -minimum that lies to the right of v . Such cases were however ruled out above, where the number of vertices v for which a terminal event on $\delta_r^{(3)}$ can be reached in this fashion was bounded by $O(k^3 n \lambda_5(n))$ in equation (1).

(a') $\eta' \in \delta_r^{(1)}$ and we make a *left* turn from $\delta_{\beta'}$ to $\delta_r^{(3)}$ at η : See Figure 13(a'). This case can only arise when $\delta_{\beta'}$ lands back on $\delta_r^{(1)}$, between v and δ_{β} , as in Figure 13(a'). (Otherwise, the component of $\tau_{r,g,y}$ that contains $\delta_r^{(3)}$ would have been forced to be bounded: It has to be contained in the region bounded by δ_{β} , $\delta_{\beta'}$, $\delta_r^{(3)}$, and $\delta_r^{(1)}$; see Figure 13(a''). This possibility has however already been ruled out, as in case (a) above.) Observe that there are overall at most $4k$ arcs $\delta_{\beta'}$ that land back on $\delta_r^{(1)}$ in this fashion. Therefore, at least $2\ell - 4k$ of the ℓ -shallow arcs emanating from $\bar{\delta}_r^{(3)}$ do




Figure 13: Various cases of ℓ -shallow blue arcs that emanate from $\bar{\delta}_r^{(3)}$.

not belong to case (a').

(b) $\eta' \in \delta_r^{(2)}$: See Figure 13(b). This case can arise only when we make a left turn from δ_β to $\delta_r^{(3)}$ at η , or else $\delta_r^{(3)}$ would have to be a bounded component, as in case (a') (the configuration would have looked like a 'mirror image' of the one depicted in Figure 13(a'')). Any red k -shallow arc that emanates from $\delta_b^{(1)}$ must then cross either δ_β or $\delta_{\beta'}$. At most k of these red arcs can cross δ_β , so at least k of them cross each ℓ -shallow arc $\delta_{\beta'}$ that falls into case (b). Since any of these k -shallow red arcs can cross only k blue arcs, it follows that at most k of the ℓ -shallow arcs emanating from $\bar{\delta}_r^{(3)}$ belong to case (b). Since only at most $4k$ arcs fall into case (a'), we conclude that one of the cases (a), (c), or (d), must arise for at least $2\ell - 5k > \ell$ arcs $\delta_{\beta'}$.

(c) η' lies to the left of v . This case cannot arise when we make a right turn from δ_β to $\delta_r^{(3)}$ at η (see Figure 13(c)), for then we could reach, in the opposite direction, a local x -extremum on $\delta_r^{(3)}$, as in case (a). However, if we make a left turn at η , as shown in Figure 13(c'), then $\delta_{\beta'}$ must leave $\delta_r^{(3)}$ in the positive x -direction, or else it would have been 'trapped' between $\delta_b^{(1)}$ on one side and $\delta_r^{(1)}$, δ_β , and $\delta_r^{(3)}$ on the other side, which would make it impossible for $\delta_{\beta'}$ to reach to the left of v . Hence, $\delta_{\beta'}$ must have a locally x -maximal point before it reaches η' ; since this is a terminal event, this contradicts the shallowness of $\delta_{\beta'}$.

(d) η' lies to the right of v , on a new edge $\delta_r^{(4)}$ of ∂W_0 , different from $\delta_r^{(i)}$, for $i = 1, 2, 3$: See Figure 13(d). In this case we trace $\delta_r^{(4)}$ from η' in the positive x -direction, and we will not reach any vertex of W_0 (we have already exhausted all such vertices to the right of v). The tracing will thus reach a terminal event. Since η' lies at combined level at most $5k + 5\ell \leq 6\ell$, this also bounds the 3-colored level of the terminal event. Hence, arguing as above, the number of such events is $O(\ell^2 n \lambda_5(n))$, and each of them is charged by only $O(\ell)$ vertices v . To see the latter claim, we spell out, for the sake of completeness, a modified version of a previous argument:

Trace $\delta_r^{(4)}$ back (in either direction, if more than one direction is applicable, as there may be two choices for r and for $\delta_r^{(4)}$) from the terminal event, until the first time we reach a point whose combined level (including the blue level) is at most 6ℓ (this will be either at η' or earlier). Then any charging vertex v is a vertex incident to $\tilde{f}_r, \tilde{f}_g, \tilde{f}_y$, and to some \tilde{f}_b , from the at most 6ℓ blue surfaces that lie below the stopping point. In other words, the terminal event can be charged by at most $O(\ell)$ vertices v , implying that the number of vertices v that fall into this final subcase is $O(\ell^3 n \lambda_5(n))$.

This completes the consideration of all possible situations that arise with vertices v of index at most 3. Collecting all of the bounds obtained during the analysis of such vertices leads to the following equation:

$$V_0^{(3)}(n) = O\left((k^2 \ell^2 + \ell^3) n \lambda_5(n) + V_0^{(2)}(n) + k^4 V_0^{(2)}\left(\frac{n}{4k}\right) + k^2 V_0^{(3)}\left(\frac{n}{5k}\right) + \ell^4 V_0^{(2)}\left(\frac{n}{5\ell}\right) + k^2 \ell^2 V_0^{(3)}\left(\frac{n}{6\ell}\right) \right),$$

which proves the lemma. \square

3.3 Putting It All Together

Recall that in Section 3.2 we have only handled regular vertices v . To complete the counting, we have to add the number of irregular vertices to each of the above bounds on the quantities $V_0^{(j)}(n)$. Since there are only $O(n\lambda_5(n))$ irregular vertices, this does not affect any of these asymptotic estimates. Thus, collecting the bounds obtained in Lemmas 3.3–3.5, we obtain the following recurrence relations:

$$\begin{aligned} V_0^{(0)}(n) &= O(n\lambda_5(n)). \\ V_0^{(1)}(n) &= O(n\lambda_5(n)). \\ V_0^{(2)}(n) &= O\left(k^3 n\lambda_5(n) + k^2 V_0^{(3)}\left(\frac{n}{5k}\right)\right). \\ V_0^{(3)}(n) &= O\left((k^2\ell^2 + \ell^3) n\lambda_5(n) + V_0^{(2)}(n) + k^4 V_0^{(2)}\left(\frac{n}{4k}\right) + \right. \\ &\quad \left. k^2 V_0^{(3)}\left(\frac{n}{5k}\right) + \ell^4 V_0^{(2)}\left(\frac{n}{5\ell}\right) + k^2 \ell^2 V_0^{(3)}\left(\frac{n}{6\ell}\right)\right). \end{aligned}$$

We choose different values of k in the recurrences for $V_0^{(2)}$ and for $V_0^{(3)}$, and denote them by k_2 and k_3 , respectively. These values, together with ℓ , are chosen to be sufficiently large constants, satisfying $\ell = k_3^{1/(c\varepsilon)}$ and $k_2 = \ell^{1/(c\varepsilon)}$, for an arbitrarily small but prescribed positive constant ε , and for some fixed small positive fraction c . (Note that this choice of parameters satisfies $\ell \gg k_3$, which was needed in our analysis.) We also require that k_3^ε be sufficiently large. The recurrence for $V_0^{(3)}$ then becomes

$$\begin{aligned} V_0^{(3)}(n) &= O\left(\left(k_2^{c\varepsilon(2+2c\varepsilon)} + k_2^{3c\varepsilon}\right) n\lambda_5(n) + V_0^{(2)}(n) + k_2^{4c^2\varepsilon^2} V_0^{(2)}\left(\frac{n}{4k_2^{c^2\varepsilon^2}}\right) + \right. \\ &\quad \left. k_2^{2c^2\varepsilon^2} V_0^{(3)}\left(\frac{n}{5k_2^{c^2\varepsilon^2}}\right) + k_2^{4c\varepsilon} V_0^{(2)}\left(\frac{n}{5k_2^{c\varepsilon}}\right) + k_2^{2c\varepsilon+2c^2\varepsilon^2} V_0^{(3)}\left(\frac{n}{6k_2^{c\varepsilon}}\right)\right) \\ &= O\left(k_2^{3+2c\varepsilon} n\lambda_5(n) + k_2^2 V_0^{(3)}\left(\frac{n}{5k_2}\right) + k_2^{2+4c^2\varepsilon^2} V_0^{(3)}\left(\frac{n}{20k_2^{1+c^2\varepsilon^2}}\right) + \right. \\ &\quad \left. k_2^{2c^2\varepsilon^2} V_0^{(3)}\left(\frac{n}{5k_2^{c^2\varepsilon^2}}\right) + k_2^{2+4c\varepsilon} V_0^{(3)}\left(\frac{n}{25k_2^{1+c\varepsilon}}\right) + k_2^{2c\varepsilon(1+c\varepsilon)} V_0^{(3)}\left(\frac{n}{6k_2^{c\varepsilon}}\right)\right). \end{aligned}$$

As in other works where similar recurrences have been derived (see, e.g., [24]), it is easy to show, using induction on n , that, with an appropriate choice of c and k_2 (where the choice of k_2 depends on ε but the choice of c does not), the solution of this recurrence is

$$V_0(n) = V_0^{(3)}(n) = O(n^{2+\varepsilon}),$$

for any ε , where the constant of proportionality depends on ε . We have thus shown

Theorem 3.6. *The complexity of the Euclidean Voronoi diagram of a set of n lines in \mathbb{R}^3 with four distinct orientations is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Remark. Inspecting the proof of Theorem 3.6, we see that it is fairly general, and does not use explicitly the fact that the sites are lines. It can thus be extended to the case of the Voronoi diagram of any reasonable collection of sites (of constant description complexity), which is the union of four subfamilies, under any reasonable metric in \mathbb{R}^3 , provided that (i) we have a near-quadratic bound for the complexity of the diagram of any three of the given families, and (ii) any four sites determine at most 8 Voronoi vertices. We strongly suspect that the requirement (ii) can be dropped. This would require us to handle vertices v that have index $x \geq 4$, which in turn would have made the preceding analysis more complicated, mainly by having to use additional thresholds for shallowness (like the k and ℓ that we used). Still, it seems plausible that the analysis could go through.

4 More than Four Orientations

The case of an arbitrary number c of orientations is easy to handle, by noting that any vertex v of the full Voronoi diagram $Vor(L)$ is also a vertex of the diagram of the set of all lines whose orientations are equal to the (at most) four orientations of the lines that are (equally) nearest to v . Let u_1, \dots, u_c denote the given orientations. Let L_j , for $j = 1, \dots, c$, denote the set of lines in L at orientation u_j , and put $n_j = |L_j|$. Then $\sum_{j=1}^c n_j = n$. Suppose, without loss of generality, that $n_1 \leq n_2 \leq \dots \leq n_c$. The number of vertices of $Vor(L)$ is at most $\sum_{i < j < k < l} V_{ijkl}$, where V_{ijkl} is the number of vertices of $Vor(L_i \cup L_j \cup L_k \cup L_l)$. By Theorem 3.6, $V_{ijkl} = O((n_i + n_j + n_k + n_l)^{2+\varepsilon}) = O(n_l^{2+\varepsilon})$. Hence, the complexity of $Vor(L)$ is at most $O\left(\sum_{i < j < k < l} n_l^{2+\varepsilon}\right) = O\left(\sum_{l=4}^c l^3 n_l^{2+\varepsilon}\right)$. As is easily verified, the maximum value of this latter sum is $O(c^3 n^{2+\varepsilon})$. We thus obtain the main result of the paper:

Theorem 4.1. *The combinatorial complexity of the Euclidean Voronoi diagram of n lines in three dimensions, where the lines have $1 \leq c \leq n$ distinct orientations, is $O(c^3 n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Corollary 4.2. *The combinatorial complexity of the Euclidean Voronoi diagram of n lines in \mathbb{R}^3 that have a constant number of distinct orientations is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Remark. As shown in [23], the complexity of the diagram in the general case, without any restrictions on the orientations of the lines (that is, when $c = n$), is $O(n^{3+\varepsilon})$. This leads us to conjecture that the bound in Theorem 4.1 can be improved to at least $O(cn^{2+\varepsilon})$, for any $\varepsilon > 0$. The latter bound is consistent with the result of [23] (when $c = O(n)$) and with Corollary 4.2 (when $c = O(1)$), and might be easier to obtain than a near-quadratic bound like $O(n^{2+\varepsilon})$ for any $1 \leq c \leq n$. (Nevertheless, in line with the general conjecture concerning 3-dimensional Voronoi diagrams, we conjecture that the latter bound does indeed hold independently of c .)

Acknowledgements

We are grateful to Mark de Berg, who raised the possibility of improving the bound $O(c^4n^{2+\epsilon})$, which appeared in preliminary versions of this paper, to the bound $O(c^3n^{2+\epsilon})$ that is currently proved in Section 4. We are also grateful to Emo Welzl and the European Graduate Program on Combinatorics, Geometry and Computation, for hosting us in ETH Zürich, where the work on this paper was initiated.

References

- [1] P. K. Agarwal and M. Sharir. Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. *Discrete and Computational Geometry*, 24:645–685, 2000.
- [2] B. Aronov. A lower bound on Voronoi diagram complexity. *Information Processing Letters*, 83:183–185, 2002.
- [3] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16:78–96, 1987.
- [4] F. Aurenhammer and R. Klein. Voronoi diagrams. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
- [5] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete and Computational Geometry*, 19(4):473–484, 1998.
- [6] L. P. Chew, K. Kedem, M. Sharir, B. Tagansky, and E. Welzl. Voronoi diagrams of lines in 3-space under polyhedral convex distance functions. *Journal of Algorithms*, 29(2):238–255, 1998.
- [7] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete and Computational Geometry*, 4:387–421, 1989.
- [8] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete and Computational Geometry*, 1:25–44, 1986.
- [9] G. Elber and M.-S. Kim. The bisector surface of freeform rational space curves. In *Proc. 13th ACM Symposium on Computational Geometry*, pages 473–474, 1997.
- [10] S. Fortune. Voronoi diagrams and Delaunay triangulations. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 20, pages 377–388. CRC Press LLC, Boca Raton, FL, 1997.

- [11] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete and Computational Geometry*, 12:313–326, 1994.
- [12] D. Halperin and M. Sharir. Almost tight upper bounds for the single cell and zone problems in three dimensions. *Discrete and Computational Geometry*, 14:385–410, 1995.
- [13] A. Harnack. Über die Vielfaltigkeit der ebenen algebraischen Kurven. *Math. Ann.*, 10:189–199, 1876.
- [14] R. Hartshorne. *Algebraic Geometry*. Springer-Verlag, New York, 1977.
- [15] D. P. Huttenlocher, K. Kedem, and J. M. Kleinberg. On dynamic Voronoi diagrams and the minimum Hausdorff distance for point sets under Euclidean motion in the plane. In *Proc. 8th ACM Symposium on Computational Geometry*, pages 110–119, 1992.
- [16] C. Icking and L. Ma. A tight bound for the complexity of Voronoi diagrams under polyhedral convex distance functions in 3D. In *Proc. 33rd ACM Symposium on the Theory of Computing*, pages 316–321, 2001.
- [17] V. Klee. On the complexity of d -dimensional Voronoi diagrams. *Archiv der Mathematik*, 34:75–80, 1980.
- [18] V. Koltun and M. Sharir. Polyhedral Voronoi diagrams of polyhedra in three dimensions. In *Proc. 18th ACM Symposium on Computational Geometry*, 2002.
- [19] D. Leven and M. Sharir. Intersection and proximity problems and Voronoi diagrams. In J. T. Schwartz and C.-K. Yap, editors, *Advances in Robotics 1: Algorithmic and Geometric Aspects of Robotics*, pages 187–228. Lawrence Erlbaum Associates, Hillsdale, NJ, 1987.
- [20] J. S. B. Mitchell and J. O’Rourke. Computational geometry column 42. *International Journal of Computational Geometry and Applications*, 2001. Also in SIGACT News 32(3):63-72, Issue 120, 2001, and on <http://www.cs.smith.edu/~orourke/TOPP>.
- [21] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. John Wiley & Sons, 2000.
- [22] R. Seidel. Exact upper bounds for the number of faces in d -dimensional Voronoi diagrams. In P. Gritzman and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 517–530. AMS Press, Providence, RI, 1991.
- [23] M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete and Computational Geometry*, 12:327–345, 1994.

- [24] M. Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- [25] B. Tagansky. *The Complexity of Substructures in Arrangements of Surfaces*. Ph.D. thesis, Tel Aviv University, Tel Aviv, 1996.

Appendix

In this appendix we provide a study of the geometric structure of *bisectors* and *trisectors*, which are, respectively, the loci of points equidistant from two and three lines. This analysis is useful in its own right, but most of the details are not needed for the main result of this paper.

We begin with the analysis of bisectors, which have also been studied, e.g., in [9]. Consider the bisector H_{ℓ_1, ℓ_2} between two lines ℓ_1, ℓ_2 at different orientations. Without loss of generality, by translating, rotating, and scaling 3-space, we may assume that ℓ_1 and ℓ_2 are both horizontal, ℓ_1 (resp., ℓ_2) passes through $(0, 0, 1)$ (resp., $(0, 0, -1)$) and forms a horizontal angle of α (resp., $-\alpha$) with the positive x -direction, for $\alpha \in (-\pi/2, \pi/2]$.

The squared distance of a point (x, y, z) from ℓ_1 is

$$d^2((x, y, z), \ell_1) = x^2 + y^2 + (z - 1)^2 - (x \cos \alpha + y \sin \alpha)^2,$$

and the squared distance of (x, y, z) from ℓ_2 is

$$d^2((x, y, z), \ell_2) = x^2 + y^2 + (z + 1)^2 - (x \cos \alpha - y \sin \alpha)^2.$$

Hence, the equation of H_{ℓ_1, ℓ_2} is

$$x^2 + y^2 + (z - 1)^2 - (x \cos \alpha + y \sin \alpha)^2 = x^2 + y^2 + (z + 1)^2 - (x \cos \alpha - y \sin \alpha)^2,$$

or

$$z = -xy \sin \alpha \cos \alpha.$$

This is the equation of a hyperbolic paraboloid. It has two sets of generating lines, one set consisting of lines parallel to the xz -plane, and the other consisting of lines parallel to the yz -plane. Specifically, lines in the first family have the following form, parametrized over $t \in \mathbb{R}$:

$$\lambda_t : \quad y = -\frac{t}{\sin \alpha \cos \alpha}, \quad z = tx.$$

Similarly, lines in the second family have the form, parametrized over $s \in \mathbb{R}$,

$$\bar{\lambda}_s : \quad x = -\frac{s}{\sin \alpha \cos \alpha}, \quad z = sy.$$

We can project H_{ℓ_1, ℓ_2} onto the xy -plane π_0 , bijectively, and note that the generating lines project to lines parallel to the axes.

Fix a line λ_t of the first family, having parameter t . Let ℓ_3 be a differently-oriented line, passing through some point $\mathbf{a} = (a_1, a_2, a_3)$ and having direction $\mathbf{u} = (u_1, u_2, u_3)$, which is a unit vector along ℓ_3 , and is common to all input lines of a fixed color. By our general position assumption, we may assume that $u_3 \neq 0$, i.e., that the direction \mathbf{u} is not coplanar with the directions of ℓ_1 and ℓ_2 . Without loss of generality, we assume that $\mathbf{a} \cdot \mathbf{u} = 0$. The distance between a point $\mathbf{w} = \mathbf{w}(x)$ on λ_t , parametrized as $(x, -t/(\sin \alpha \cos \alpha), tx)$, and ℓ_3 is

$$d^2(\mathbf{w}, \ell_3) = \|\mathbf{w} - \mathbf{a}\|^2 - ((\mathbf{w} - \mathbf{a}) \cdot \mathbf{u})^2 = \|\mathbf{w} - \mathbf{a}\|^2 - (\mathbf{w} \cdot \mathbf{u})^2 = \\ (x - a_1)^2 + \left(\frac{t}{\sin \alpha \cos \alpha} + a_2\right)^2 + (tx - a_3)^2 - \left(xu_1 - \frac{tu_2}{\sin \alpha \cos \alpha} + txu_3\right)^2.$$

Consider the function

$$F(x) = d^2(\mathbf{w}(x), \ell_3) - d^2(\mathbf{w}(x), \ell_1) = \\ (x - a_1)^2 + \left(\frac{t}{\sin \alpha \cos \alpha} + a_2\right)^2 + (tx - a_3)^2 - \left(xu_1 - \frac{tu_2}{\sin \alpha \cos \alpha} + txu_3\right)^2 \\ - x^2 - \frac{t^2}{\sin^2 \alpha \cos^2 \alpha} - (tx - 1)^2 + \left(x \cos \alpha - \frac{t}{\cos \alpha}\right)^2.$$

Note that $F(x)$ is positive (resp., zero, negative) if the ball centered at $\mathbf{w}(x)$ and touching ℓ_1 and ℓ_2 is disjoint from (resp., touches, intersects) ℓ_3 . Hence, the locus of the roots of $F(x)$, as we trace them by varying t from $-\infty$ to $+\infty$, is the *trisector* $\tau_{\ell_1, \ell_2, \ell_3}$ —the locus of all centers of balls that touch ℓ_1, ℓ_2, ℓ_3 simultaneously.

The function $F(x)$ is quadratic (for any fixed t), and its global behavior along λ_t depends largely on the sign of the coefficient of x^2 , which is

$$A(t) = 1 + t^2 - (u_1 + tu_3)^2 - 1 - t^2 + \cos^2 \alpha = \cos^2 \alpha - (u_1 + tu_3)^2.$$

Hence, if $A(t) > 0$ then $F(x)$ is convex, and is positive at $x = \pm\infty$, meaning that at the extremities of λ_t the ball touching ℓ_1, ℓ_2 is disjoint from ℓ_3 (we are in the *free region* associated with ℓ_3), whereas if $A(t) < 0$ then $F(x)$ is concave, and at the extremities of λ_t we are in the *intersection region* of ℓ_3 .

In other words, assuming, as above, that $u_3 \neq 0$, and, for specificity, that $u_3 > 0$, we have that $A(t) < 0$ iff $|u_1 + tu_3| > \cos \alpha$, or

$$t > \frac{-u_1 + \cos \alpha}{u_3} \quad \text{or} \quad t < \frac{-u_1 - \cos \alpha}{u_3}.$$

The corresponding critical y -values are

$$y_T = \frac{u_1 + \cos \alpha}{u_3 \sin \alpha \cos \alpha} \quad \text{and} \quad y_B = \frac{u_1 - \cos \alpha}{u_3 \sin \alpha \cos \alpha},$$

and we denote the corresponding horizontal *critical lines* by $\lambda^{(T)}$, $\lambda^{(B)}$, respectively. (Note that the critical lines depend *only* on the orientation \mathbf{u} of ℓ_3 .)

We next apply a symmetric analysis to lines in the other family. We obtain that the critical x -values where the corresponding quadratic function changes from being convex to being concave are

$$x_R = \frac{u_2 + \sin \alpha}{u_3 \sin \alpha \cos \alpha} \quad \text{and} \quad x_L = \frac{u_2 - \sin \alpha}{u_3 \sin \alpha \cos \alpha};$$

the corresponding vertical critical lines are denoted by $\bar{\lambda}^{(R)}$, $\bar{\lambda}^{(L)}$.

We next claim that, for $|t|$ sufficiently large, the line λ_t intersects the trisector in exactly two points. For this, we need to show that the discriminant of the quadratic equation $F(x)$ becomes positive as $|t|$ tends to ∞ .

Write $F(x)$ as $A(t)x^2 + 2B(t)x + C(t)$, where

$$\begin{aligned} A(t) &= \cos^2 \alpha - (u_1 + tu_3)^2 \\ B(t) &= -a_1 - a_3 t + \frac{u_2 t (u_1 + tu_3)}{\sin \alpha \cos \alpha} \\ C(t) &= a_1^2 + \left(\frac{t}{\sin \alpha \cos \alpha} + a_2 \right)^2 + a_3^2 - \frac{t^2 u_2^2}{\sin^2 \alpha \cos^2 \alpha} - \frac{t^2}{\sin^2 \alpha \cos^2 \alpha} - 1 + \frac{t^2}{\cos^2 \alpha}. \end{aligned}$$

As $|t|$ tends to ∞ , the sign of the discriminant $\Delta(t)$ depends only on the coefficients of t^2 in these three expressions. That is, the limit of Δ/t^4 is

$$\begin{aligned} \lim_{|t| \rightarrow \infty} \frac{B^2(t) - A(t)C(t)}{t^4} &= \lim_{|t| \rightarrow \infty} \left[\left(\frac{B(t)}{t^2} \right)^2 - \frac{A(t)}{t^2} \cdot \frac{C(t)}{t^2} \right] = \\ &= \frac{u_2^2 u_3^2}{\sin^2 \alpha \cos^2 \alpha} + u_3^2 \cdot \left(\frac{1}{\cos^2 \alpha} - \frac{u_2^2}{\sin^2 \alpha \cos^2 \alpha} \right) = \frac{u_3^2}{\cos^2 \alpha} > 0. \end{aligned}$$

That is, for large values of $|t|$, the trisector $\tau_{\ell_1, \ell_2, \ell_3}$ meets λ_t at two points $w_1(t)$, $w_2(t)$. The asymptotic values of these roots are:

$$\begin{aligned} \lim_{|t| \rightarrow \infty} w_{1,2}(t) &= \lim_{|t| \rightarrow \infty} \frac{-B(t) \pm \sqrt{\Delta(t)}}{A(t)} = \lim_{|t| \rightarrow \infty} \frac{-B(t)/t^2 \pm \sqrt{\Delta(t)/t^4}}{A(t)/t^2} = \\ &= \frac{-\frac{u_2 u_3}{\sin \alpha \cos \alpha} \pm \frac{u_3}{\cos \alpha}}{-u_3^2} = \frac{u_2 \pm \sin \alpha}{u_3 \sin \alpha \cos \alpha}. \end{aligned}$$

That is, $w_1(t)$ and $w_2(t)$ tend to x_L and x_R , respectively.

Symmetrically, there always exist two intersection points of $\tau_{\ell_1, \ell_2, \ell_3}$ with the lines $\bar{\lambda}_s$, as $|s|$ tends to ∞ , and their limits are at the ordinates y_B and y_T .

We have thus shown that any sufficiently large circle intersects the trisector at 8 points. We denote the points ‘at infinity’ that lie on the vertical critical lines $\bar{\lambda}^{(L)}$, $\bar{\lambda}^{(R)}$, as $v_{LB}, v_{LT}, v_{RB}, v_{RT}$, where v_{LB} (resp., v_{LT}) is the bottom (resp., top) end of $\bar{\lambda}^{(L)}$, and similarly for the other two points. The points at infinity on the horizontal lines are denoted, in a similar manner, as $h_{LB}, h_{LT}, h_{RB}, h_{RT}$. See Figure 14(a) for an illustration.

Assuming that the trisector is non-singular, it has exactly 4 unbounded components, each connecting two of these points at infinity. We next proceed to classify the structure of these components.




Figure 14: The various possible structures of a trisector.

The function $F(x)$ becomes linear along each of the horizontal critical lines $\lambda^{(T)}$, $\lambda^{(B)}$, and thus each of these two critical lines is intersected by the trisector exactly once; symmetrically, this also holds for $\bar{\lambda}^{(L)}$, $\bar{\lambda}^{(R)}$. Number the 8 points at infinity in a cyclic order. Then it is clear that each odd-numbered point must be connected to an even-numbered point, since the components of the trisector are disjoint. Hence, v_{LT} can be connected either to h_{LT} , or to v_{LB} , or to h_{RB} , or to v_{RT} , and similarly for the other points at infinity.

Consider the second case, in which v_{LT} is connected to v_{LB} via one component γ_1 of the trisector. This component crosses the two critical horizontal lines $\lambda^{(B)}$, $\lambda^{(T)}$ (each exactly once). In this case, no other component of the trisector can intersect any of these lines, so each of the remaining three components is fully contained in one of the three horizontal slabs delimited by $\lambda^{(B)}$ and $\lambda^{(T)}$, and each of these slabs contains exactly one such component. It then follows that these components must connect h_{LT} to h_{LB} , v_{RT} to h_{RT} , and v_{RB} to h_{RB} . Moreover, γ_1 must cross $\bar{\lambda}^{(L)}$ (exactly once), and one of the two components on the right must cross $\bar{\lambda}^{(R)}$ (exactly once). Hence, the trisector has a shape similar to that shown in Figure 14(b).

Consider next the third case, in which v_{LT} is connected to h_{RB} via one component γ_1 of the trisector. Another component, γ_2 , must connect v_{RT} to h_{RT} . We have two subcases:

In the first subcase, h_{LT} is connected to v_{RB} and h_{LB} is connected to v_{LB} . In this case, none of the components can cross any of its asymptotes. See Figure 14(a).

In the second subcase, h_{LT} is connected to h_{LB} and v_{LB} is connected to v_{RB} . In this case, we must allow each of the lines $\lambda^{(B)}$, $\bar{\lambda}^{(L)}$ to be crossed (once) by some component. See Figure 14(c). This figure depicts one of several possible subcases, depending on which component crosses which critical line. In Figure 14(c) the component connecting v_{LT} to h_{RB} crosses all four critical lines, but it might also be possible for this component to cross only $\lambda^{(T)}$ and $\bar{\lambda}^{(R)}$ or to cross just one more critical line, and let the left and/or bottom components cross the other one or two critical lines (in a manner similar to that of the top-right component in Figure 14(b)).

If none of the above cases occur, including their various symmetric variants, then each end of each critical line must be connected to one of its two neighbors in the above cyclic order. Only two cases are possible:

In the first subcase, h_{LT} is connected to h_{LB} , v_{LB} is connected to v_{RB} , h_{RT} is connected to h_{RB} , and v_{LT} is connected to v_{RT} . As above, we must let some of these components cross some of their asymptotes, to ensure that each of the four critical lines is crossed once by the trisector. See Figure 14(d), which, as above, depicts just one of several possible subcases.

In the second subcase, h_{LT} is connected to v_{LT} , h_{LB} is connected to v_{LB} , h_{RT} is connected to v_{RT} , and h_{RB} is connected to v_{RB} . Again, we must let some of these components cross some of their asymptotes. One of several possible such configurations is shown in Figure 14(e).

We also note that each trisector is an algebraic curve of degree 4. By Harnack's theorem [13], the number of components of a real nonsingular algebraic plane curve of degree d is at most $(d - 1)(d - 2)/2 + 1$. Hence the number of components of each trisector is at most $3 \cdot 2/2 + 1 = 4$. Since it has exactly four unbounded components, we conclude that these are *all* the components of the trisector. In particular, no component of any trisector is bounded. This completes the classification of the trisectors.

Remark. We conjecture that, up to symmetry, only trisectors of the kind shown in Figure 14(b) are possible. A program that we have written to explore the structure of trisectors has only revealed trisectors of this kind, after several tens of thousands of tests with randomly generated lines.