Voronoi Diagrams of Lines in 3-Space Under Polyhedral Convex Distance Functions*

L. Paul Chew† Klara Kedem‡ Micha Sharir§

Boaz Tagansky¶ Emo Welzl‖

Abstract

The combinatorial complexity of the Voronoi diagram of n lines in three dimensions under a convex distance function induced by a polytope with a constant number of edges is shown to be $O(n^2\alpha(n)\log n)$, where α is a slowly growing inverse of the Ackermann function. There are arrangements of n lines where this complexity can be as large as $\Omega(n^2\alpha(n))$.

1 Introduction

Statement of result. Let P be a closed convex polytope in 3-space which contains the origin. Given any point $w \in \mathbb{R}^3$, its distance from a line (or any other object) ℓ, as induced by P, is

$$d_P(w,\ell) = \inf\{t \geq 0 \,:\, (w+tP)\cap \ell \neq \emptyset\} \;;$$

*Work by Paul Chew and Klara Kedem has been supported by AFOSR Grant AFOSR-91-0328 and by the U.S.–Israeli Binational Science Foundation. Work by Paul Chew has also been supported by ONR Grant N00014-89-J-1946 and ARPA under ONR contract N00014-88-K-0591. Work by Micha Sharir and Emo Welzl has been supported by the G.I.F. — the German Israeli Foundation for Scientific Research and Development (work by Boaz Tagansky has also been supported by this grant), and by a Max Planck Research Award. Work by Micha Sharir has also been supported by National Science Foundation Grants CCR-91-22103 and CCR-93-11127, and by grants from the U.S.–Israeli Binational Science Foundation, and the Israel Science Fund administered by the Israeli Academy of Sciences. Work by Emo Welzl has also been supported by the EC Basic Research Action, Project ALCOM II.

†Department of Computer Science, Cornell University, Ithaca, NY 14853, USA, `chew@cs.cornell.edu`

‡Department of Mathematics and Computer Science, Ben Gurion University, Beer-Sheva 84105, Israel, `klara@ivory.bgu.ac.il`

§School of Mathematical Sciences, Tel Aviv University, and Courant Institute of Mathematical Sciences, New York University, `sharir@math.tau.ac.il`

¶School of Mathematical Sciences, Tel Aviv University, `mia@math.tau.ac.il`

‖Institut für Informatik, Freie Universität Berlin, Takustr. 9, D-14195 Berlin, Germany, `emo@inf.fu-berlin.de`

d_P is called the *convex distance function* induced by P. (Note that the L_1- and L_∞-metrics are special cases of this distance function, obtained by taking P to be an octahedron or a cube, respectively.)

Let $\mathcal{L} = \{\ell_1, \ldots, \ell_n\}$ be a collection of n lines in three dimensions. The *Voronoi diagram* $\mathrm{Vor}_P(\mathcal{L})$ of $\mathcal{L}$ induced by P is defined as the decomposition of 3-space into *Voronoi cells*, one cell for each line ℓ_i in $\mathcal{L}$, defined by

$$V(\ell_i) = \{w \,:\, d_P(w,\ell_i) \le d_P(w,\ell_j) \text{ for all } j \ne i\} \,.$$

Each Voronoi cell is a polyhedron (in general, not convex). The *combinatorial complexity* of $\mathrm{Vor}_P(\mathcal{L})$ is the overall number of faces (vertices, edges, facets) of the diagram. See [2] for a comprehensive survey of known results (mostly in two dimensions) on Voronoi diagrams.

We show that the combinatorial complexity of $\mathrm{Vor}_P(\mathcal{L})$ is $O(n^2\alpha(n)\log n)$, where the constant of proportionality depends on the number of edges of P. This is an improvement over previous bounds, which are super-cubic in n: (1) a bound of $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, proved in [10] for more general Voronoi diagrams in three dimensions; and (2) a bound of $O(n^3\alpha(n))$, for the case we consider here, which can be developed by using the results of [8]. There are arrangements of lines where $\mathrm{Vor}_P(\mathcal{L})$ has complexity $\Omega(n^2\alpha(n))$, even if P is a tetrahedron. Thus there remains a gap of a logarithmic factor between the lower and upper bound.

A prevailing conjecture is that the complexity of fairly general three-dimensional Voronoi diagrams is near-quadratic in the number of sites. However, quadratic or near-quadratic bounds for this complexity were known only in the special case of point sites and the Euclidean distance [6] (see also [3], where techniques from this paper have recently been applied to the case of point sites and the L_1-metric, and [4], where an $O(n^2\alpha(n))$ bound is obtained for the case of lines in 3-space and a distance function induced by a flat (i.e. two-dimensional) convex *polygon*). Our proof crucially depends on P being a polytope, so it does not seem to extend to the case of Voronoi diagrams of lines under the Euclidean distance (where the best upper bound known on the complexity of the diagram is still $O(n^{3+\varepsilon})$ [10]).

Generalized Voronoi diagrams, as defined above, are strongly related to the union of Minkowski sums of polyhedra as studied in [1]. Specifically, consider a generalized setup, in which $\mathcal{L}$ is a set of pairwise-disjoint polyhedral sites $A_1, \ldots, A_n$. Then the boundary of the union $U_t = \bigcup_{i=1}^n A_i \oplus (-tP)$ is the locus of all points whose smallest d_P distance to any site is t. Thus the union of Minkowski sums can be regarded as a 'cross-section' of the Voronoi diagram of the sites A_i under the distance function d_P. That is, if we interpret the Voronoi diagram as the (projection of the) lower envelope in $\mathbb{R}^4$ of the n distance functions to the sites (as in [7]), then ∂U_t is the projection of the cross-section of the lower envelope with a horizontal hyperplane at height t. Using fairly complicated topological analysis, the paper [1] establishes a near-quadratic bound on the complexity of this union. The problem that we study in this paper is thus considerably more difficult, in the sense that we bound the complexity of the entire diagram – not just of a cross section (however, [1] can handle arbitrary polyhedral sites instead of lines).

The approach. We assume that $\mathcal{L}$ and P are in general position (we will be more specific about that below). The proof bounds the number of *Voronoi vertices*. There are three classes of Voronoi vertices: The main class consists of points v that are equidistant (under d_P) to four distinct lines in $\mathcal{L}$, and this distance is the smallest from v to any line in $\mathcal{L}$. If we denote this distance by ρ, then v corresponds to a homothetic placement $\overline{P} = \overline{P}(v, \rho) = v + \rho P$ that touches four lines in $\mathcal{L}$ and no line in $\mathcal{L}$ intersects the interior of $\overline{P}$. The other two possible classes of vertices in the Voronoi diagram (and corresponding placements of homothetic copies $\overline{P}$ of P) are: (2) two lines in $\mathcal{L}$ touch $\overline{P}$ at edges and one line in $\mathcal{L}$ touches $\overline{P}$ at a vertex, and (3) two lines in $\mathcal{L}$ touch $\overline{P}$ at vertices. Vertices of types (2) and (3) are relatively easy to count. It is also easy to show that it suffices to bound the number of vertices, since the number of all other faces of the diagram is proportional to the number of vertices, if one assumed general position of the sites (see below). Even when this assumpotion is not made, the bound that we will derive will bound the overall complexity of the diagram.

In Section 2 we assume that P is a tetrahedron denoted by Δ. We will repeatedly consider the motion of a homothetic copy $\overline{\Delta}$ of Δ where three given lines keep contact with three given edges of $\overline{\Delta}$, while the tetrahedron is allowed to expand or shrink and translate. If no line intersects the interior of the tetrahedron during this motion, it corresponds to moving along an edge of the Voronoi diagram. However, in our analysis we will also be moving in an 'opposite' direction: Starting from a *free* homothetic copy $\overline{\Delta}$ with four line-edge contacts and no line intersecting the interior (as in (1) above), we 'slide' Δ while maintaining three of the contacts, but forcing the fourth line to enter the interior of Δ. In this process we look for the next *critical placement*, where lines in $\mathcal{L}$ meet edges or vertices of Δ. We charge the original placement to other placements reached by performing this sliding process in two different ways (by maintaining different triples of contacts). This leads to a recurrence relation, which we can solve by using a probabilistic argument, adapted from a technique recently introduced in [12].

Section 3 extends the result to general polytopes P. The proof there is more complicated, and proceeds by induction on the number of vertices of P. We first argue that it suffices to consider polytopes P with up to 8 vertices. We classify the Voronoi vertices into several classes, depending on the pattern of contacts that occur at the placement corresponding to the vertex, and analyze each class separately. These additional analyses are simpler than the one used in the case of a tetrahedron.

A few concluding remarks and open problems are given in Section 4.

2 The Case of a Tetrahedron

We assume in this section that P is a tetrahedron, denoted as Δ, and that $\mathcal{L}$ is a fixed set of n lines in space. A homothetic copy $\overline{\Delta} = \overline{\Delta}(z, \rho) = z + \rho\Delta$, for $z \in \mathbb{R}^3$, $\rho > 0$, is called a *placement* of Δ. If f is a face (vertex, edge, or facet) of Δ then $\overline{f}$ refers to the corresponding face of placement $\overline{\Delta}$. If a line $\ell \in \mathcal{L}$ intersects an edge or

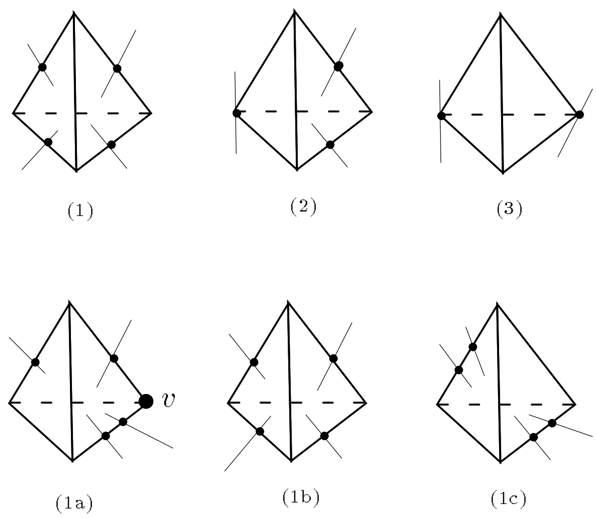

Figure 1: Cases of rigid contacts; in (1a), v is a vertex 'incident' to three contacts

a vertex $\overline{f}$ of $\overline{\Delta}$ then we call the pair (f, ℓ) a *contact* of $\overline{\Delta}$; it is called a *vertex contact* or an *edge contact*, depending on whether f is a vertex or an edge, respectively. The contact is called *touching* if ℓ does not intersect the interior of $\overline{\Delta}$. A placement $\overline{\Delta}$ is *free* if no line in $\mathcal{L}$ intersects the interior of $\overline{\Delta}$, and is called *almost free* if only (some of) those lines involved in contacts pierce the interior of $\overline{\Delta}$. A placement is called *rigid*, if either (1) there are four edge contacts, or (2) there are two edge contacts and one vertex contact, or (3) there are two vertex contacts (see Figure 1).

We assume that Δ and $\mathcal{L}$ are in *general position* in the following sense: (i) No two lines in $\mathcal{L}$ have a common point or are parallel. (ii) No line in $\mathcal{L}$ is parallel to any facet of Δ. (iii) No line parallel to an edge of Δ can touch three lines of $\mathcal{L}$. (iv) No placement can have more contacts than those prescribed for rigid placements in cases (1), (2) and (3) of the preceding paragraph. (v) If two rigid placements have the same contacts, then they are identical.[1] These assumptions can be enforced by an infinitesimal perturbation of the lines in $\mathcal{L}$. These assumptions involve no real loss of generality, because, as can be shown, the maximum complexity of $\mathrm{Vor}_\Delta(\mathcal{L})$, for a set $\mathcal{L}$ of n lines in 3-space, is obtained when $\mathcal{L}$ and Δ are in general position. We extend the notion of general position to arbitrary convex polytopes P, by imposing similar conditions on the lines in $\mathcal{L}$. Moreover, we require that the same conditions should also hold with repsect to any subpolytope of P, defined as the convex hull of a subset of the vertices of P. The property that the maximum complexity of the diagram is obtained when the lines and P are in general position hold here as well.

The placements with four edge contacts are further discriminated (see Figure 1), depending on whether (1a) at least three of the edge contacts occur on (not neces-

[1] If $\overline{\Delta}(z, \rho)$ is rigid then, as is easily checked, the four parameters z, ρ satisfy a system of four linear equations. Condition (v) requires that this system always have a unique solution.

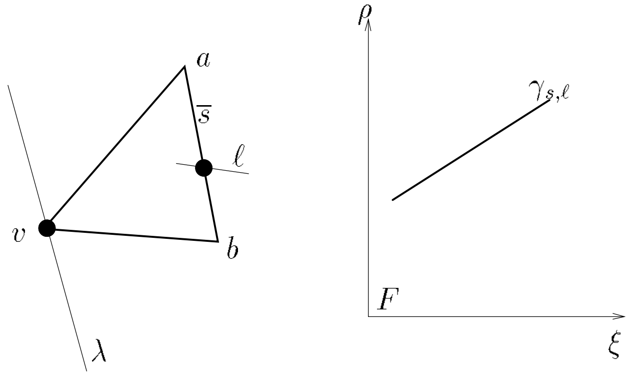

Figure 2: The vertex v moves along λ while $\overline{s}$ contacts ℓ (left), and the corresponding locus $\gamma_{s,\ell}$ in the frame F (right)

sarily distinct) edges incident to a common vertex of Δ (*common-vertex-contacts*), or (1b) there are four edges involved in contacts, which form a quadrilateral in space (*quadrilateral-contacts*), or (1c) there are two edges involved in contacts, which do not share a common vertex, and each of them has two line contacts (*opposite-edge-contacts*). These three cases classify *all* possibilities of four edge contacts.

Our goal is to bound the number of free rigid placements. We first settle the easy cases (1a), (2), (3), and then deal with the more involved situations in (1b) and (1c).

Rigid placements with vertex contacts and a lower bound. Because of the general position assumption, two prescribed vertex contacts can be obtained by at most one placement. That is, there are at most $12\binom{n}{2}$ rigid placements with two vertex contacts (not even requiring that these placements be free).

Next consider a triangle T (still in three dimensions) with vertex v, and a line λ (not necessarily in $\mathcal{L}$). All homothetic placements $\overline{T}$ of T with $\overline{v}$ on λ can be represented in a two-dimensional frame F, parametrized by (ξ, ρ), where ξ represents the position of v on λ and ρ is the scaling factor of $\overline{T}$ (see Figure 2). Let s be the edge of T opposite to v, and let ℓ be a line in $\mathcal{L}$. We want to show that the placements of T where $\overline{s}$ intersects ℓ can be represented by a (possibly empty) straight segment or ray $\gamma_{s,\ell}$ in the frame F.

Let a and b be the endpoints of s. If $\overline{v}$ is incident to λ, then $\overline{a}$ has to lie in a plane Π which contains λ and is parallel to the edge connecting v and a. We can parametrize the placements with contact (v, λ) by the respective positions of a in Π (each such position uniquely determines a placement $\overline{T}$; in fact, F and Π have a correspondence by an affine mapping, so let us continue the discussion in Π). Moreover, let us allow negative scaling factors for the time being (using the term scaling instead of placement). Note that one halfplane Π^+ of Π bounded by λ represents positive scalings (i.e. placements), and the other halfplane Π^- represents negative scalings.

Whenever $\overline{s}$ is in contact with line ℓ, then $\overline{a}$ is in a plane Π' that contains ℓ and is parallel to s. Thus, for scalings with contacts (v, λ) and (s, ℓ), the vertex $\overline{a}$ has to be on a line μ that is the intersection of Π and Π'. There is a point p_a on this line μ which corresponds to the scaling with contacts (v, λ) and (a, ℓ); p_a is the intersection of ℓ with Π. Similarly, there is a point p_b representing the position of a when the contacts (v, λ) and (b, ℓ) occur. The locus $\delta_{s,\ell}$ of $\overline{a}$ for scalings with contacts (v, λ) and (s, ℓ) is either given by the segment connecting p_a and p_b, (when both p_a and p_b lie on the same side of λ) or by the complement of this segment within μ, (if λ separates p_a and p_b). That is, $\delta_{s,\ell} \cap \Pi^+$ is empty, a line segment, or a ray.

The discussion should also make it clear that, given a segment γ in Π^+, there exists a line ℓ so that γ is the locus of $\overline{a}$ for placements with contacts (v, λ) and (s, ℓ); actually, there are two such lines, because we can choose either endpoint of γ to represent p_a.

We return to the frame F, and use the term 'segments' for both bounded line segments and rays. All free placements of T with v on λ are constrained by the lower envelope (in direction ρ) of all segments $\gamma_{s,\ell}$, for $\ell \in \mathcal{L}$. The lower envelope of n such segments has complexity $O(n\alpha(n))$, where α is a slowly growing inverse of the Ackermann function [11]. It is also known that there are sets of line segments which attain this bound [11, 13].

We use these observations to bound the number of placements of the tetrahedron Δ with vertex contacts, and for a lower bound of the complexity of the Voronoi diagram of lines in space. Let us show right away that there are n lines in space, so that there are $\Omega(n^2\alpha(n))$ rigid placements of a triangle T (with vertex v and opposite edge s). In the frame F as described above, we choose $\lceil n/2 \rceil$ line segments above the ξ-axis whose lower envelope has complexity $\Theta(n\alpha(n))$. Let $\mathcal{L}'$ be a set of corresponding lines in space. Now there are $\Theta(n\alpha(n))$ free placements of T, where v touches λ and s touches two lines in $\mathcal{L}'$. A small perturbation of λ will not change the complexity of the lower envelope in the frame. So we choose another set $\mathcal{L}''$ of $\lfloor n/2 \rfloor$ lines which are perturbed versions of λ. Some care is needed: we have to ensure that those lines do not meet the free rigid placements at which v lies on other lines in $\mathcal{L}''$. To this end, we let Π be a plane containing λ and parallel to s; so if $\overline{v}$ sits on Π, then Π does not intersect the interior of $\overline{T}$. Now we can choose the lines in $\mathcal{L}''$ in Π parallel to λ, but close enough so that their frames still give $\Theta(n\alpha(n))$ complexity for the respective lower envelopes. Altogether, this gives $\Omega(n^2\alpha(n))$ free rigid placements of T. If desired, a final sufficiently small perturbation will achieve general position. Of course, we can also extend the triangle T to an (almost flat) tetrahedron Δ without destroying any of the free rigid placements counted above.

We continue with the upper bound argument and bound the number of rigid placements with one vertex contact.

Let v be a vertex of Δ and let ℓ be a line in $\mathcal{L}$. As described above, we can represent all placements with vertex contact (v, ℓ) in a two-dimensional frame F. Every edge s in Δ not incident to v, and every line $\ell' \in \mathcal{L} \setminus \{\ell\}$ determines a straight segment in this frame. (If s is incident to v then, by our general position assumption, v must lie at a unique point on ℓ when the double contact (v,ℓ), (s,ℓ') occurs. In other words, the segment induced by s in F is vertical, namely, parallel to the ρ-axis. We ignore these edges since they can be shown to induce only a linear number of contacts of the type considered here.) All choices of s and ℓ' yield a two-dimensional arrangement of at most $3(n-1)$ line segments in F. The lower envelope (in direction ρ) of these segments represents all free placements of Δ with contact (v,ℓ) and with at least one extra *touching* contact. A segment endpoint that lies on this lower envelope represents a rigid placement of Δ with two vertex contacts, and an intersection of two segments on the envelope gives a rigid free placement with one vertex and two edge contacts. As stated above, the complexity of this lower envelope is $O((n-1)\alpha(n))$.

In what we counted so far, the extra contacts (beyond (v,ℓ)) must be touching. It is easy to see that any almost-free rigid placement (with the contact (v,ℓ)) must appear on a shallow level of the arrangement of these segments (namely, at level at most two, where the lower envelope is counted as level 0). The complexity of these levels is still bounded by $O((n-1)\alpha(n))$ (see, e.g., [5, 9]). There are four choices for v and n choices for ℓ, and so there are at most $O(n(n-1)\alpha(n))$ almost-free rigid placements with one vertex contact.

We say that a line in $\mathcal{L}$ *violates* a placement, if it intersects the interior of the placement, but it is not involved in a contact. Let $D_k(\mathcal{L})$ be the number of rigid placements with one vertex contact, and at most k violating lines. We have just shown that $D_0(\mathcal{L}) = O(n^2\alpha(n))$. It is also easy to show that $D_1(\mathcal{L}) = O(n^2\alpha(n))$: Using 2-dimensional representations by planes Π as above, it is easily seen that almost-free rigid placements with a vertex contact and with one violating line appear at level at most 4 in the corresponding arrangements of segments within the planes Π. Using the analysis of [5, 9] as above, the asserted bound on $D_1(\mathcal{L})$ follows.

Three contacts incident to a common vertex. Let $\overline{\Delta}$ be an almost-free rigid placement, where three of the contacts appear on (not necessarily distinct) edges incident to a common vertex $\overline{v}$ (i.e., a placement of type (1a)). If we shrink the tetrahedron while keeping $\overline{v}$ fixed, the contacts on edges incident to $\overline{v}$ will continue to exist until one of them becomes a vertex contact (u,ℓ). The fourth contact is broken by this shrinkage, but its participating line might still intersect the interior of $\overline{\Delta}$. That is, after we stop at the vertex contact, we have a rigid placement $\overline{\Delta}$ of Δ with one vertex contact (say (u,ℓ)) and two edge contacts (say (s,ℓ') and (t,ℓ'')), and possibly one extra line in $\mathcal{L}$ intersecting the interior of $\overline{\Delta}$. Note that we can expand $\overline{\Delta}$ again, maintaining contact with the same three lines, in at most two ways: either reversing the shrinking process about $\overline{v}$, if $s \neq t$ (in this case v is the unique common endpoint of s and t), or, if $s = t$, also about the other endpoint of s, until a new edge contact is created. Thus we have bounded the number of almost-free rigid placements of type (1a) by twice the number of rigid placements of type (2) where at most one line not involved in a contact pierces the interior. As noted above, this gives a bound of $O(n^2\alpha(n))$.

Let $E_k(\mathcal{L})$ be the number of rigid placements of type (1a) with at most k violating

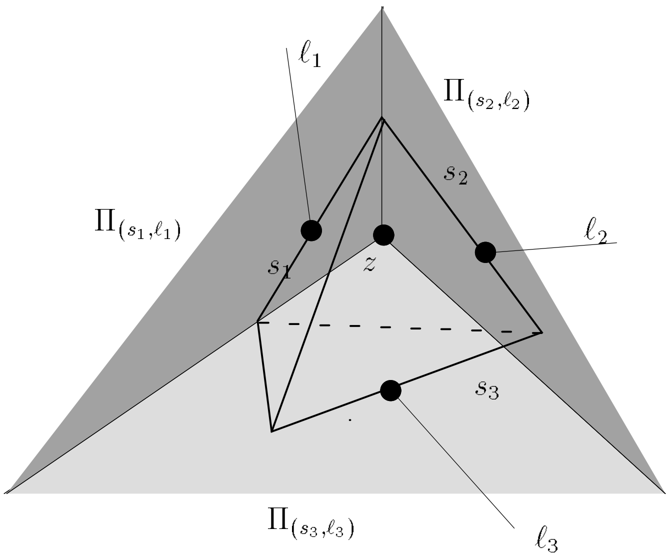

Figure 3: Sliding the tetrahedron along three contacts is achieved by scaling it with respect to z

lines. We have shown that $E_0(\mathcal{L}) = O(n^2\alpha(n))$. A straightforward application of the probabilistic arguments of [5, 9] shows that $E_1(\mathcal{L})$ is also $O(n^2\alpha(n))$.

The bounds on $D_1(\mathcal{L})$ and $E_1(\mathcal{L})$ will be used in proving bounds for the remaining patterns of contacts, namely, quadrilateral- and opposite-edge-contacts. In the following discussion, we refer to the facets incident to an edge $\overline{s}$ of a placement $\overline{\Delta}$ as *s-facets*. Clearly, every line that intersects $\overline{\Delta}$ but has no edge contact must intersect two facets, and there is a unique edge s for which these facets are the two s-facets.

Sliding a tetrahedron. A line $\ell \in \mathcal{L}$ and an edge s of Δ define a unique plane $\Pi_{(s,\ell)}$ that contains ℓ and is parallel to s. In order for the edge $\overline{s}$ of a placement $\overline{\Delta}$ to meet a line ℓ, $\overline{s}$ must lie in $\Pi_{(s,\ell)}$. Now let ℓ_1, ℓ_2, ℓ_3 be three lines in $\mathcal{L}$, and let s_1, s_2, s_3 be edges of Δ, not all three the same. The planes $\Pi_{(s_m,\ell_m)}$, for $m = 1, 2, 3$, meet at a common point z (as implied by the general position assumption, see Figure 3). All placements where $\overline{s_m} \subseteq \Pi_{(s_m,\ell_m)}$, for all $m = 1, 2, 3$, can be obtained by scaling one such placement with respect to z as a center. Every line ℓ defines a (possibly empty) interval of positive scaling factors for which ℓ meets the sliding tetrahedron.[2] If the intervals where the lines ℓ_1, ℓ_2, ℓ_3 meet, respectively, the edges s_1, s_2, s_3 have a nonempty intersection I, then we call the corresponding motion, in which Δ is scaled with respect to z as a center and the scaling factor lies in I, *"sliding Δ with contacts (s_m, ℓ_m), for $m = 1, 2, 3$"*.

We consider a rigid placement $\overline{\Delta}$ with four touching edge contacts, together with

[2] More generally, if P and C are convex bodies and z is a point, then the range of positive reals λ for which $P \cap (z + \lambda(C - z)) \neq \emptyset$ is an interval.

an *ordered* pair $((s, \ell'), (t, \ell''))$ of its contacts. The tuple $(\overline{\Delta}, (s, \ell'), (t, \ell''))$ is called a *doubly-hinged rigid placement.* Such a doubly-hinged rigid placement $\overline{\Delta}$ is called an (i, j)-placement if it has an opposite-edge- or quadrilateral-contact, if $s \neq t$, and if there are exactly $i + j$ lines intersecting the interior of $\overline{\Delta}$, so that i of them intersect the two s-facets, and j of them intersect the two t-facets. Every free rigid placement with quadrilateral-contact induces twelve $(0, 0)$-placements (there are 12 choices for the ordered pair (s, t), each with a unique choice of ℓ' and ℓ''), and every opposite-edge-contact induces eight $(0, 0)$-placements (there are two choices for (s, t), two choices for ℓ' and two for ℓ''). The number of (i, j)-placements, for a given set $\mathcal{L}$ of lines, is denoted by $C_{(i,j)}(\mathcal{L})$.

The charging scheme for quadrilateral- and opposite-edge-contacts. Our strategy is to charge every $(0, 0)$-placement to *two* other placements (with at most one violation).

Given a $(0, 0)$-placement $(\overline{\Delta}, (s, \ell'), (t, \ell''))$, we slide the tetrahedron while releasing the first contact (s, ℓ') (and maintaining the other three contacts) in the direction that causes ℓ' to penetrate the tetrahedron. (It is easily seen that there is a unique such direction.) The process is stopped as soon as a new contact occurs. We discriminate the following types of events:

(A) *A new edge contact with a line ℓ^* is encountered on edge s.* Note that we have reached a $(1, 0)$-placement $(\overline{\overline{\Delta}}, (s, \ell^*), (t, \ell''))$ (because ℓ' still intersects the two s-facets), and there is exactly one $(0, 0)$-placement (the one we started with) from which this $(1, 0)$-placement can be obtained in the prescribed manner. This can be seen by simply reversing the process, that is, we release the line ℓ^* involved in the first contact (s, ℓ^*) and slide in the unique direction which does not cause ℓ^* to penetrate the tetrahedron. When the next contact occurs, we have reached our initial $(0, 0)$-placement. (In fact, if we start with an arbitrary $(1, 0)$-placement, this reversed process need not end in a $(0, 0)$-placement! Note also that this argument, and the arguments used below, make use of our general position assumption.)

(B) *A new edge contact is encountered on an edge different from s.* We have four edge contacts, which are common-vertex-contacts (this holds, by an easy consideration, because we started with a quadrilateral-contact or with an opposite-edge-contact). Either we have four touching contacts and one violating line, or we have three touching contacts and one contact whose line penetrates the interior of the placement (this is the line ℓ'). In both cases, applying a 'reverse' sliding motion, as above, one can show that such a placement can be reached at most a constant number of times in the presecribed manner, and the number of such placements, as argued above, is $O(n^2\alpha(n))$.

(C) *One of the current edge contacts becomes a vertex contact.* Again, such a placement can be reached only in a constant number of ways from a $(0, 0)$-placement, and the number of such placements is $O(n^2\alpha(n))$. (To see this, slide Δ backwards, as above, by forcing the line incident to the vertex to touch one of the three edges that meet at that vertex.)

In a similar way we slide while releasing the second contact (t, ℓ'') in the direction that causes ℓ'' to pierce the tetrahedron. When the next contact occurs, we have either reached a $(0,1)$-placement, or we are in a situation as described in (B) and (C) above.

Now we charge every $(0,0)$-placement to both placements we have reached by sliding as above. The number of placements which are of type (B) or (C) has already been bounded by $O(n^2\alpha(n))$, so that the charging scheme gives us an inequality of

$$2C_{(0,0)}(\mathcal{L}) \le C_{(1,0)}(\mathcal{L}) + C_{(0,1)}(\mathcal{L}) + O(n^2\alpha(n)) \ ,$$

or

$$2C_0(\mathcal{L}) \le C_1(\mathcal{L}) + O(n^2\alpha(n)) \ , \tag{1}$$

if we write $C_0(\mathcal{L})$ and $C_1(\mathcal{L})$ instead of $C_{(0,0)}(\mathcal{L})$ and $C_{(1,0)}(\mathcal{L})+C_{(0,1)}(\mathcal{L})$, respectively. We next adapt the probabilistic analysis technique of [12]. That is, let $\mathcal{R}$ be a random subset of $\mathcal{L}$ of cardinality $n-1$. Then we get

$$\mathrm{E}[C_0(\mathcal{R})] \ge \frac{n-4}{n} C_0(\mathcal{L}) + \frac{1}{n} C_1(\mathcal{L}) \ , \tag{2}$$

since a $(0,0)$-placement in $\mathcal{R}$ can be derived (i) either from a $(0,0)$-placement in $\mathcal{L}$, if none of its four contact lines is omitted from $\mathcal{R}$ (which happens with probability $(n-4)/n$) or (ii) from a $(1,0)$- or $(0,1)$-placement in $\mathcal{L}$, if the violating line from $\mathcal{L}$ is omitted from $\mathcal{R}$ (which happens with probability $1/n$). The reason that we need not have an equality in (2) is that there might exist other $(0,0)$-placements $\overline{\Delta}$ in $\mathcal{R}$ where the line omitted from $\mathcal{R}$ pierces $\overline{\Delta}$ at a pair of facets that are not both s-facets and are not both t-facets.

If we substitute $C_1(\mathcal{L})$ in (1) using (2), we obtain

$$(n-2)C_0(\mathcal{L}) \le n\mathrm{E}[C_0(\mathcal{R})] + O(n^2\alpha(n)) \ . \tag{3}$$

We define $C_0(k)$ as the maximum value of $C_0(\mathcal{K})$ over all sets $\mathcal{K}$ of k lines; in particular, $C_0(1) = C_0(2) = C_0(3) = 0$. Then (3) (divided by $n(n-1)(n-2)$) implies

$$\frac{C_0(n)}{n(n-1)} \le \frac{C_0(n-1)}{(n-1)(n-2)} + O\left(\frac{\alpha(n)}{n}\right) \ ,$$

for $n \ge 4$, which immediately gives us a bound of $O(n^2\alpha(n)\log n)$ on $C_0(n)$.

We have thus obtained the main result of this section:

Theorem 2.1 *The number of free rigid placements of a tetrahedron Δ among n lines $\mathcal{L}$ in space is bounded by $O(n^2\alpha(n)\log n)$. Hence, the complexity of* $\mathrm{Vor}_\Delta(\mathcal{L})$ *is $O(n^2\alpha(n)\log n)$. There exist tetrahedra and sets of n lines, for any n, for which this number is $\Omega(n^2\alpha(n))$.* ◻

3 The General Polyhedral Case

We next extend Theorem 2.1 to convex distance functions induced by an arbitrary fixed convex polytope P. As in the previous section, we need to bound the number of free rigid placements. We first argue that it suffices to consider only polytopes with up to 8 vertices — as long as we consider the number of edges of P to be constant. Then we extend the result from 4 vertices (the tetrahedral case) to larger polytopes by induction on the number of vertices.

Lemma 3.1 *Let $r_8(n)$ bound the number of free rigid placements of a convex polytope with at most 8 vertices in an arrangement of n lines. Then $\binom{m}{4}r_8(n)$ bounds the number of free rigid placements among n lines of any convex polytope P with m edges.*

Proof. Let $\overline{P}$ be a free rigid placement of P. Let $S = \{s_1, s_2, s_3, s_4\}$ be four edges of $\overline{P}$ which cover all contacts of P; that is, every edge involved in a contact is in S, and every vertex involved in a contact is incident to some edge in S. Now $\overline{Q_S} = \text{conv}(\overline{s_1} \cup \overline{s_2} \cup \overline{s_3} \cup \overline{s_4})$ is a free rigid placement of $Q_S = \text{conv}(s_1 \cup s_2 \cup s_3 \cup s_4)$, a polytope with at most 8 vertices. As Q_S has at most $r_8(n)$ free rigid placements for each S, and there are at most $\binom{m}{4}$ choices for S, the assertion of the lemma follows. □

We now proceed by induction on the number k of vertices of P. Assume now that P has k vertices, $5 \le k \le 8$, and that we have already shown a bound of $A_{k-1}n^2\alpha(n)\log n$, for some constant A_{k-1} (depending only on k), for the number of free rigid placements of polytopes with $k-1$ or fewer vertices among n lines in space. The base case of $k = 4$ is provided by Theorem 2.1. We will show that there are at most $A_k n^2\alpha(n)\log n$ free rigid placements of P, for a larger constant A_k (that depends only on k). The proof proceeds through the following steps:

Lemma 3.2 *The number of free rigid placements of P with a vertex contact, or with three edge contacts on (not necessarily distinct) edges with a common vertex, is bounded by $O(n^2\alpha(n))$.*

Proof. This is shown exactly as in the case of a tetrahedron. □

A notation is required for the following induction steps. For a polytope P and a vertex x of P, let P_x denote the convex hull of all vertices of P except for x.

Lemma 3.3 *Let a be a vertex of P. The number of free rigid placements, where vertex a is neither involved in a contact, nor it is incident to an edge involved in a contact, is bounded by $O(n^2\alpha(n)\log n)$.*

Proof. Such a placement corresponds to a rigid placement of P_a which has $k-1$ vertices. Hence, by the induction hypothesis, the number of such rigid placements is bounded as claimed. (Clearly, here k must be at least 5.) □

Lemma 3.4 *Let s be an edge of P. There are at most $O(n^2\alpha(n)\log n)$ free rigid placements of P with four edge contacts where there is exactly one contact on edge s, and no other contact edge shares a vertex with s.*

Proof. Let a and b be the vertices incident to s. Given a free rigid placement $\overline{P}$ as described in the lemma, start sliding P from $\overline{P}$ so that the line ℓ with contact on s pierces P, and the remaining three contacts are retained. During this process, we follow the corresponding motions of P_a and P_b. Note that both polytopes share the three contacts with P, and that they are free in the initial part of the motion, since ℓ will not intersect them. We stop the motion as soon as we reach either a vertex contact on P_a or P_b, or a new edge contact appears on P_a or P_b. By the induction hypothesis, the number of such terminal placements is $O(n^2\alpha(n)\log n)$. The process can be reversed only in a constant number of ways. Here it is important to observe that, as we stop the sliding, line ℓ still intersects the interior of $\overline{P}$. Indeed, in order for this line to escape P, it has to sweep through s or through one of the polytopes P_a or P_b. It cannot meet s, because that is where we came from, and as soon as ℓ touches P_a or P_b, we stop. □

For the remaining situations, all involving free rigid placements with four edge contacts, it is helpful to consider the *contact graph* G induced by those contacts. Its nodes are all the vertices of the polytope that are incident to edges with contacts, and for each contact on an edge of P we connect the two incident vertices by an arc in the graph G. That is, if there are two contacts on edge s, then its incident vertices are connected by *two* arcs in the graph (recall that three contacts or more on an edge are excluded by the general position assumption). G will always have four arcs, and at most 8 nodes. Lemma 3.2 deals with the situation when there is a node of degree three or four in this graph, and Lemma 3.4 addresses the case where there is an arc with both incident nodes having degree one in the contact graph.

What is left? Contact graphs with all nodes having degree one or two, and no connected component formed by a single arc. That is, there are at most two connected components which are either cycles or paths of length at least two. This leaves us with the following possibilities for G (see Figure 4): (i) G is a cycle of length 4. (ii) G consists of two disjoint cycles of length 2. (iii) G is a path of length 4. (iv) G is composed of a cycle of length 2 and a path of length 2. (v) G is formed by two paths of length 2 each.

In cases (i) and (ii), the contact graph has four nodes, and we are either facing the base case of a tetrahedron, or there is vertex of the polyhedron which is incident to no contact edge, in which case Lemma 3.3 comes into play.

Cases (iii)–(v) have the following property in common: We can enumerate the arcs of G as e_1, e_2, e_3, e_4, so that e_1 and e_2 have a common node c, e_3 is incident to a node a of degree one in G, and e_4 is incident to a node b of degree one in G; e_1 and e_2 may share both incident nodes, (i.e., may be induced by the same edge of P,) but otherwise the arcs in G stem from distinct edges of P. This pattern is treated next:

Lemma 3.5 *Let s_1, s_2, s_3, s_4 be edges of P which are mutually distinct, except that s_1 and s_2* may *be equal. Moreover, suppose that s_1 and s_2 share a common vertex*

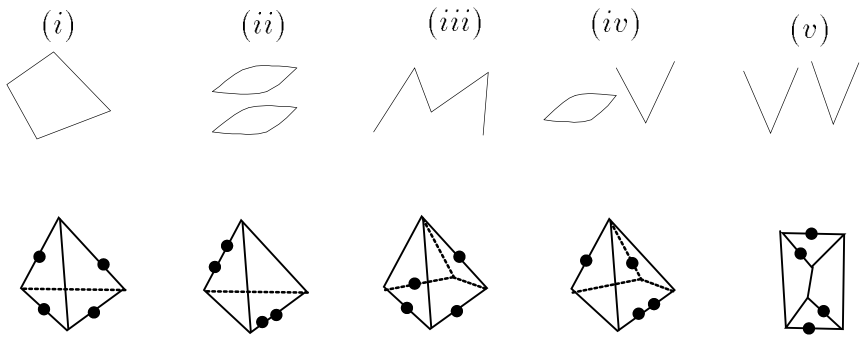

Figure 4: Various contact graphs (top) and rigid placements that realize them (bottom)

c, that there is a vertex a incident to s_3 only (among s_1, s_2, s_3, s_4), and that there is a vertex b incident to s_4 only. Then there are at most $O(n^2\alpha(n)\log n)$ free rigid placements of P with four edge contacts on s_1, s_2, s_3, and s_4.

Proof. Note that both s_1 and s_2 are present in P_a and in P_b, that s_3 is present in P_b, and that s_4 is present in P_a.

Fix lines ℓ' and ℓ''. In all placements of P with edge contacts (s_1, ℓ') and (s_2, ℓ''), the vertex $\bar{c}$ lies on a line λ, the intersection of planes $\Pi_{(s_1,\ell')}$ and $\Pi_{(s_2,\ell'')}$. Now fix a coordinate frame $F = F_{\ell',\ell''}$ representing all such placements. We can parametrize F by the two parameters (ξ, ρ), where ξ gives the position of c along λ, and ρ is the scaling factor of P. As described in the previous section, every edge of P not incident to c and every line in $\mathcal{L}$ defines a straight segment in F. The free placements with c on λ are represented in F above the ξ-axis and below the lower envelope E of these segments.

In the same frame F we can also represent the free placements of P_a and P_b with edge contacts (s_1, ℓ') and (s_2, ℓ''). The lower envelope E for P has to lie below the lower envelopes E_a and E_b for P_a and P_b, respectively. In particular, a vertex in E which corresponds to a free rigid placement of P with edge contacts on s_3 and s_4, is a vertex in the lower envelope E' of the envelopes E_a and E_b. Since E_a and E_b are piecewise-linear functions, the complexity of E' is linear in the complexity of E_a and E_b (see, e.g., [11]). It remains to observe that, by the induction hypothesis, the complexity of all the E_a's, over all pairs of lines ℓ' and ℓ'' in $\mathcal{L}$, is bounded by $O(n^2\alpha(n)\log n)$ (and similar for b instead of a). ◻

This concludes the inductive argument from $k-1$ to k, since every pattern of contacts of a free rigid placement induces one of the situations covered by Lemmas 3.2–3.5. Hence we have established the bound on the number of free rigid placements, which, at last, entails the main result of this paper.

Theorem 3.6 *The complexity of the Voronoi diagram of a set of n lines in 3-space, under a convex distance function induced by a convex polytope P, is $O(n^2\alpha(n)\log n)$. The constant of proportionality depends on P and is cm^4, for an absolute constant c, where m the number of edges of P. There are sets of n lines, for any n, for which the Voronoi diagram has complexity $\Omega(n^2\alpha(n))$, even if P is a triangle or a tetrahedron.* ◻

Recall that the L_1- and L_∞-metrics are convex distance funtions induced by an octahedron and a cube, respectively. Hence:

Corollary 3.7 *The complexity of the Voronoi diagram of a set of n lines in 3-space under the L_1-metric or under the L_∞-metric is bounded by $O(n^2\alpha(n)\log n)$.* ◻

Remark. It is noteworthy that in the above analysis the bound of $O(n^2\alpha(n)\log n)$ comes directly from Theorem 2.1 (with the exception of the lower-order term $O(n^2\alpha(n))$ in Lemma 3.2). Thus any improvement in the analysis of the tetrahedral case would immediately yield a similar improvement for general polyhedra.

REMARK. Of course, we could continue the induction on the number of vertices of P beyond $k = 8$, but this would introduce an exponential dependence on k. This is the reason for switching to the argument in Lemma 3.1.

4 Conclusion

In this paper we have obtained the first (sub-cubic and) near-quadratic bound on the complexity of generalized Voronoi diagrams in 3-space.

There are many open problems raised by the results of this paper. First, there still remains a logarithmic gap between the upper and lower bounds that we proved. We have recently extended our proof to obtain an $O(n^2\alpha(n)\log n)$ upper bound for the complexity of the Voronoi diagram of n *line segments* under a *tetrahedral* distance function. For this scenario, we also have an $\Omega(n^2(\alpha(n))^2)$ lower bound. Arbitrary polyhedral sites and polyhedral distance functions seem to be more difficult. One reason is that there are additional types of contact (between vertices of sites and faces of P and between faces of sites and vertices of P), which make the analysis considerably more involved. Even the case of an arbitrary polyhedral distance function and point sites is still open (no sub-cubic upper bound is known). A partial progress on this problem was recently made in [3], where a tight worst-case bound of $\Theta(n^2)$ was obtained for the Voronoi diagram of n points in $\mathbb{R}^3$ under the L_1-metric. (This paper also provides tight worst-case bounds for the complexity of Voronoi diagrams of point sites in any dimension, under the L_∞-metric.)

The real challenge is, however, to extend our results to the case of Voronoi diagrams under the Euclidean distance (where P is a ball). Our proof technique relies crucially on the fact that P is polyhedral, and at present we do not see any way to extend the technique to the Euclidean case. Equally challenging is the problem of bounding the complexity of planar dynamic Voronoi diagrams under the Euclidean distance, especially the case where the sites are points, each moving along some straight line at some constant velocity (each site has its own line and velocity). This problem can be transformed into the problem of analyzing the complexity of $\text{Vor}_D(\mathcal{L})$ in three dimensions, where D is a horizontal disk and $\mathcal{L}$ is a collection of n lines in space. Again, our technique fails in this 'non-polyhedral' case.

Acknowledgments. We wish to thank Boris Aronov, Bernard Chazelle, and Jirka Matoušek for useful discussions concerning the problems studied in this paper.

References

[1] B. Aronov and M. Sharir, On translational motion planning in three dimensions, *Proc. 10th ACM Symp. on Computational Geometry* (1994), 21–30. (Also to appear in *SIAM J. Comput.*)

[2] F. Aurenhammer, Voronoi diagrams—A survey of a fundamental geometric data structure, *ACM Computing Surveys* 23 (1991), 346–405.

[3] J.-D. Boissonnat, M. Sharir, B. Tagansky and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Proc. 11th ACM Symp. on Computational Geometry* (1995), 79–88.

[4] L.P. Chew, Near-quadratic bounds for the L_1 Voronoi diagram of moving points, *Proc. 5th Canadian Conf. on Computational Geometry* (1993), 364–369.

[5] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[6] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer Verlag, Heidelberg, 1987.

[7] H. Edelsbrunner and R. Seidel, Voronoi diagrams and arrangements, *Discrete Comput. Geom.* 1 (1986), 25–44.

[8] J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, *Discrete Comput. Geom.* 4 (1989), 291–309.

[9] M. Sharir, On k–sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* 6 (1991), 593–613.

[10] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[11] M. Sharir and P. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[12] B. Tagansky, A new technique for analyzing substructures in arrangements, *Proc. 11th ACM Symp. on Computational Geometry* (1995), 200–210. (Also to appear in *Discrete Comput. Geom.*)

[13] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport–Schinzel sequences by segments, *Discrete Comput. Geom.* 3 (1988), 15–47.