Polyhedral Voronoi Diagrams of Polyhedra in Three Dimensions*

Vladlen Koltun† Micha Sharir‡

ABSTRACT

We show that that the complexity of the Voronoi diagram of a collection of disjoint polyhedra in 3-space that have n vertices overall, under a convex distance function induced by a polyhedron with $O(1)$ facets, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. We also show that when the sites are n segments in 3-space, this complexity is $O(n^2\alpha(n)\log n)$. This generalizes previous results by Chew et al. [9] and by Aronov and Sharir [4], and solves an open problem put forward by Agarwal and Sharir [2]. Specific distance functions for which our results hold are the L_1 and the L_∞ metrics. These results imply that we can preprocess a collection of polyhedra as above into a near-quadratic data structure that can answer δ-approximate Euclidean nearest-neighbor queries amidst the polyhedra in time $O(\log(n/\delta))$, for an arbitrarily small $\delta > 0$.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical algorithms and problems—Geometrical problems and computations

General Terms

Theory, Algorithms.

†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

*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.

SoCG'02, June 5-7, 2002, Barcelona, Spain.

Keywords

Voronoi diagrams, Three dimensions, Polyhedral sites, Polyhedral distance function.

1. INTRODUCTION

The Voronoi diagram of a set Γ of objects ('sites') in some space under some metric (or distance function) 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 from the combinatorial and algorithmic points of view has a long and rich history in computational geometry, beginning with the very papers that launched this field in the 1970s [17]. Following the intensive research conducted since, properties of Voronoi diagrams in the plane are very well understood, with respect to many different distance functions and types of sites (see, e.g., [6, 11] for recent comprehensive surveys of the subject).

In higher dimensions, however, some very basic problems concerning Voronoi diagrams are still wide open, and have withstood repeated attacks. One such problem is determining the combinatorial complexity of the Voronoi diagram of a set of n 'simply-shaped' sites in 3-space under a simple metric, where the prevailing conjecture is that this complexity is near-quadratic, as suggested by the best-known lower bound. However, the best upper bound derived so far is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$ [19], and even this bound is non-trivial to obtain; it is a corollary to the result that states that the lower envelope of an arrangement of semi-algebraic functions of constant description complexity in $\mathbb{R}^{d+1}$ has complexity $O(n^{d+\varepsilon})$, for any $\varepsilon > 0$ [18]. Since the Voronoi diagram of semi-algebraic sites under a semi-algebraic metric (all of constant description complexity) in d dimensions can be represented as such an envelope, as observed by Edelsbrunner and Seidel [10], this result implies that the complexity of such Voronoi diagrams is $O(n^{d+\varepsilon})$, for any $\varepsilon > 0$, as well.

A special case of this conjecture, which is still open, is the case of the Voronoi diagram of pairwise-disjoint polyhedral sites with a total of n edges in 3-space under the Euclidean metric. Two results with quite involved proofs lend credence to it. One result, due to Agarwal and Sharir [3], shows that the locus of points in $\mathbb{R}^3$ that lie at distance exactly α from the closest site has complexity $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. The second result, recently obtained by the authors [16], shows that the Voronoi diagram of n lines in 3-space under the Euclidean metric has complexity $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, if the lines have a constant number of distinct orientations.

A different research avenue is to consider Voronoi diagrams under a 'polyhedral distance function' induced by a convex polytope with a constant number of facets (see Section 2.1 for details). It is this direction that the current paper takes, and we refer to such diagrams as *polyhedral Voronoi diagrams*. The polyhedral distance functions include the well-known L_1 and L_∞ metrics, and are also interesting due to the fact that the Euclidean ball can be approximated arbitrarily well by a convex polytope. This implies that any Euclidean Voronoi diagram can be approximated with an arbitrarily high degree of accuracy by a polyhedral one. The results presented in this paper concerning polyhedral Voronoi diagrams, as well as some results that were known beforehand, are markedly better than the parallel known results for the Euclidean case.

A tight worst-case bound of $\Theta(n^2)$ has recently been presented by Icking and Ma [14] for the complexity of a polyhedral Voronoi diagram of points in $\mathbb{R}^3$. This followed earlier works by Tagansky [20], who has derived a bound of $O(n^2 \log n)$ for this complexity (and a worst-case bound of $\Theta(n^2)$ for the L_1-metric), and by Boissonnat et al. [7] who have showed the $\Theta(n^2)$ bound for some special cases, and have also given a tight worst-case bound of $\Theta(n^{\lceil d/2 \rceil})$ for the case of points in d dimensions under the L_∞ metric or a distance function induced by a simplex.

Perhaps more significantly, in light of the state of the art in the Euclidean case, an upper bound of $O(n^2\alpha(n)\log n)$ for the polyhedral Voronoi diagram of n *lines* in $\mathbb{R}^3$ was proved by Chew et al. [9], together with a lower bound of $\Omega(n^2\alpha(n))$. This followed an earlier work of Chew [8] who showed a bound of $O(n^2\alpha(n))$ for the Voronoi diagram of lines in $\mathbb{R}^3$ under a distance function induced by a 2-dimensional polygon.

In conclusion to their paper [9], Chew et al. have put forward the problem of obtaining a near-quadratic upper bound for the complexity of the polyhedral Voronoi diagram of line segments, and, more generally, of polygons and polyhedra in three dimensions. It has since been restated by Agarwal and Sharir in their survey [2, Open Problem 6(ii)].[1]

In this paper we settle this problem by proving a bound of $O(n^2\alpha(n)\log n)$ for the polyhedral Voronoi diagram of n line segments in 3-space (Section 4), and a bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, for the polyhedral Voronoi diagram of a collection of disjoint polyhedra in 3-space with n vertices altogether (Section 5). (The constant of proportionality is quartic in the number of facets of the polytope that defines the distance function.) This also significantly generalizes the result of Aronov and Sharir [4] (see also [5]), who used fairly complicated topological arguments to show a near-quadratic bound for the complexity of the locus of points at any fixed (polyhedral) distance α from their nearest site.

Our results can be applied to show (see Section 6) that a collection of disjoint polyhedra in 3-space with n vertices altogether can be preprocessed into a data structure of size $O(n^{2+\varepsilon}/\delta^4)$, for any $\varepsilon > 0$, such that this data structure can answer δ-approximate Euclidean nearest-neighbor queries amidst the polyhedra in time $O(\log(n/\delta))$, for an arbitrarily small $\delta > 0$. (That is, the query returns a site whose distance to the query point is at most $1 + \delta$ times the distance to the nearest site.) To our knowledge, no such data structure with comparable performance was available before. For the case of point sites, a near-linear approximate nearest-neighbor data structure has recently been presented by Har-Peled [13].

Some of the basic techniques we employ are inspired by ideas introduced by Chew et al. [9]. We extensively utilize the probabilistic analysis method developed by Tagansky [20], commonly known as the *Tagansky technique*. We also rely on the technique of counting schemes, originally introduced by Halperin and Sharir [12, 18], and refined in [1, 15] (see [19] for more details concerning this technique).

[1] We are aware of a subsequent study by Chew of the case of segment sites and a distance function induced by a tetrahedron; as far as we know, this work has not been published.

2. PRELIMINARIES

2.1 Definitions and a Reduction to Triangles

Let $\mathcal{P}$ be a convex polytope in $\mathbb{R}^3$ with a constant number of vertices, such that $\mathcal{P}$ contains the origin in its interior. We will refer to the origin as the *center* of $\mathcal{P}$. The distance function induced by $\mathcal{P}$ is denoted by $d_\mathcal{P}$, and the distance from any point $v \in \mathbb{R}^3$ to a (possibly infinite) set of points $S \subset \mathbb{R}^3$ under $d_\mathcal{P}$ is

$$d_\mathcal{P}(v, S) = inf\{t \geq 0 : (v + t\mathcal{P}) \cap S \neq \emptyset\}.$$

The distance function $d_\mathcal{P}$ is a metric if $\mathcal{P}$ is centrally symmetric with respect to the origin.

The Voronoi diagram $Vor_\mathcal{P}(\Gamma)$ of a set Γ of m disjoint sites in 3-space is the subdivision of $\mathbb{R}^3$ into m cells, one cell for each site of Γ, such that the cell $V(\gamma)$, for $\gamma \in \Gamma$, is

$$V(\gamma) = \{v : d_\mathcal{P}(v, \gamma) \leq d_\mathcal{P}(v, \gamma'),\ \forall \gamma' \neq \gamma\}.$$

If Γ is a set of points, segments, or piecewise linear surfaces, then each $V(\gamma)$ is a (not necessarily convex) polyhedron. The vertices (edges, faces) of all $V(\gamma)$, for $\gamma \in \Gamma$, are the vertices (resp., the edges, the faces) of $Vor_\mathcal{P}(\Gamma)$. The *combinatorial complexity* of $Vor_\mathcal{P}(\Gamma)$ is the number of faces of all dimensions of $Vor_\mathcal{P}(\Gamma)$.

Let Γ' be a collection of disjoint closed polyhedra in 3-space with n vertices altogether. Throughout this paper, we assume that Γ' is in general position with respect to the polytope $\mathcal{P}$ that induces the distance function. That is, no two vertices in the scene lie on a line that is parallel to one of the faces of $\mathcal{P}$, no homothetic copy of $\mathcal{P}$ touches more than four sites of Γ' with its boundary (while otherwise being disjoint from these sites), no line parallel to an edge of $\mathcal{P}$ intersects more than two edges of Γ', etc.

This assumption is essential, as the complexity of $Vor_\mathcal{P}(\Gamma)$ can reach $\Omega(n^3)$ when the sites are in a degenerate configuration [7, Theorem 7.1].

By triangulating $\mathcal{P}$ and the boundaries of the polyhedral sites, and by applying an infinitesimal perturbation to $\mathcal{P}$ and the sites, we may assume that all faces of $\mathcal{P}$ are triangles, and that the sites consist of $O(n)$ pairwise disjoint triangles in general position.[2]

It is easy to see that $Vor_\mathcal{P}(\Gamma)$ does not contain edges and facets that are not adjacent to a vertex. The complexity of $Vor_\mathcal{P}(\Gamma)$ is thus proportional to the number of its vertices, and it is therefore sufficient to provide a bound on this quantity.

[2] This replacement of solid sites by their bounding triangles can increase the complexity of the diagram by partitioning the points in the interior of a site among the Voronoi cells of its boundary triangles.

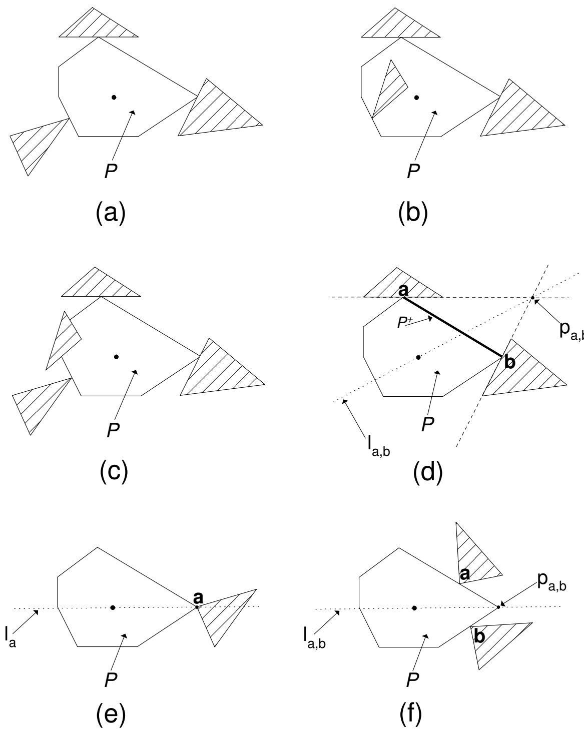

Figure 1: An illustration of some of the definitions. Details are given in Section 2.5.

2.2 Vertices of the Diagram

Some of the concepts introduced in the remainder of this section are illustrated in Figure 1.

Placements. Let Γ be a collection of triangles in general position in $\mathbb{R}^3$. Consider a homothetic copy in $\mathbb{R}^3$ of $\mathcal{P}$. It has four degrees of freedom, three for the location of its center and one for its scale relative to the center. (That is, we represent a homothetic copy $\underline{x} + \lambda\mathcal{P}$ by the quadruple $(\underline{x}, \lambda)$.) Any set of four parameters that specifies location and scale as above is called a *placement* of $\mathcal{P}$. A placement is said to be *free* if the corresponding copy of $\mathcal{P}$ is disjoint from all the sites of Γ in its interior.

Uni-contacts. A placement can be such that $\partial\mathcal{P}$ touches a triangle of Γ. This can happen in several ways. We say that a *uni-contact* (sometimes referred to as just a 'contact') occurs if a vertex of $\mathcal{P}$ touches the interior of the triangle (a V-contact), or if the relative interior of an edge of $\mathcal{P}$ touches the relative interior of an edge of the triangle (an E-contact), or if the interior of a face of $\mathcal{P}$ touches a vertex of the triangle (an F-contact). Instead of saying that a contact *occurs*, we will sometimes say that $\mathcal{P}$ *makes* or *maintains* this contact. Notice that if we force $\mathcal{P}$ to maintain a certain contact, it retains three degrees of freedom. For instance, fixing the location of the center of $\mathcal{P}$ and forcing $\mathcal{P}$ to make a certain contact uniquely determines the scale of $\mathcal{P}$.

Multi-contacts. A *bi-contact* is said to occur if a vertex of $\mathcal{P}$ touches an edge of some triangle of Γ (a V^2-bi-contact), or if an edge of $\mathcal{P}$ touches a vertex of the triangle (an E^2-bi-contact). Forcing $\mathcal{P}$ to maintain a bi-contact leaves it with only two degrees of freedom. A *tri-contact* is said to occur if a vertex of $\mathcal{P}$ touches a vertex of a triangle, and we denote this tri-contact by V^3. Maintaining such a tri-contact leaves $\mathcal{P}$ with only one degree of freedom — its center can only move along a line (more accurately, a half-line) in $\mathbb{R}^3$.

Events. Certain placements of $\mathcal{P}$ can make more than one contact. An *event* is a placement of $\mathcal{P}$ in which it makes such contacts that together leave it with zero degrees of freedom. That is, any infinitesimal movement of $\mathcal{P}$ in any direction away from an event necessarily results in losing one of the contacts that are involved in the event. An event is said to be *free* if it is a free placement. An event is said to be a uni-contact event if $\mathcal{P}$ only makes uni-contacts in this placement, and it is said to be a multi-contact event otherwise.

OBSERVATION. There is a one-to-one correspondence between vertices of $Vor_{\mathcal{P}}(\Gamma)$ and free events of $\mathcal{P}$ among Γ.

Freedom. A contact is said to be *clean* if the triangle of Γ that is incident to $\partial\mathcal{P}$ in this contact is disjoint from the interior of $\mathcal{P}$. Otherwise, the contact is said to be *dirty.* An event is said to be pseudo-free if one or more of the involved contacts are dirty, but no triangle that is not involved in a contact intersects $\mathcal{P}$. A 1-level free (resp., pseudo-free) event is an event that would be free (resp., pseudo-free) had exactly one triangle been removed from Γ. An event is said to be a 0-level event if it is either free or pseudo-free.

Degrees. Consider a placement of $\mathcal{P}$ that is an event. The *degree* of a vertex v of $\mathcal{P}$ in this event is defined as the number of distinct contacts that involve either v or an edge or a face of $\mathcal{P}$ that are incident to v. The degree of an edge e of $\mathcal{P}$ is defined as the number of contacts that involve the relative interiors of e or of the two faces of $\mathcal{P}$ that are incident to e. The degree of a face f of $\mathcal{P}$ is defined as the number of contacts that involve the interior of f.

V-, V^2- and V^3-contacts increase the degree of one vertex; E- and E^2-contacts increase the degree of an edge and the two vertices adjacent to it; F-contacts increase the degree of a face, along with the degrees of the three edges and the three vertices adjacent to it. Notice that the general position assumption implies that the degree of a face is at most 1. The degree of an edge or a vertex can reach 4, which is the maximum possible degree, as is again implied by the general position assumption.

Incidence. Consider a placement of $\mathcal{P}$ that is an event. Two contacts α and β are said to be incident in one vertex (resp., two vertices) in this event if the feature of $\mathcal{P}$ that is involved in α shares one vertex (resp., two vertices) with the feature of $\mathcal{P}$ that is involved in β.

Activeness. Let i be the number of vertices of $\mathcal{P}$. We say that features of $\mathcal{P}$ that have strictly positive degree in some event are 'active' in that event. The general position assumption implies that not more that four distinct contacts can participate in an event. Since each contact increases only the degrees of the vertices that are adjacent to the feature of $\mathcal{P}$ that is involved in this contact, and each feature (face, edge, or vertex) is adjacent to at most three vertices, at most 12 vertices can be active in an event. This implies that it suffices to consider polytopes $\mathcal{P}$ with 12 or fewer vertices. Indeed, if $\mathcal{P}$ has more than 12 vertices then any event of $\mathcal{P}$ is also an event of a sub-polytope $\mathcal{P}'$ that has 12 or fewer vertices that are adjacent to 4 features of $\mathcal{P}$ (clearly, if a placement of $\mathcal{P}$ is free, then the corresponding placement of $\mathcal{P}'$ is also free.)

The maximal number of vertices of $Vor_{\mathcal{P}}(\Gamma)$, for a polytope $\mathcal{P}$ with i vertices, is thus proportional to the maximal number of vertices of $Vor_{\mathcal{Q}}(\Gamma)$, where $\mathcal{Q}$ has 12 or fewer vertices. Moreover, the dependence on i in the constant of proportionality is quartic, since it suffices to consider all sub-polytopes of $\mathcal{P}$ that are defined by 4 or fewer distinct features of $\mathcal{P}$. The number of sub-polytopes with at most 12 vertices that have to be considered is therefore less than i^4. Hence, $N_i = O(i^4 N_{12})$, for all $i > 12$, where $N_i = N_i(\Gamma)$ denotes the maximum number of vertices of $Vor_{\mathcal{P}}(\Gamma)$, for a given Γ, when $\mathcal{P}$ has i vertices.

Moreover, in order to bound the maximal number of events of a certain combinatorial type, it suffices to consider only polytopes $\mathcal{P}$ that have a number of vertices that is equal to the maximum possible number of active vertices in events of this type. For example, it is enough to consider tetrahedral $\mathcal{P}$ (four vertices) to bound the maximal number of events that involve four V-contacts.

2.3 Sliding Along Three Contacts

A basic paradigm that will prove very useful in our analysis is that of 'sliding' along three contacts.

Consider an E-contact α that involves an edge e of some triangle of Γ and an edge e' of $\mathcal{P}$. Let Π_α be the plane that contains e and is parallel to e'. 'Sliding' $\mathcal{P}$ along α means translating it in some direction, and possibly also scaling it, such that the contact α is maintained throughout the movement. It is obvious that during such sliding, the edge e' of $\mathcal{P}$ that is involved in α has to move inside the plane Π_α.

If α is an F-contact, involving a vertex v of some triangle of Γ and a face f of $\mathcal{P}$, sliding is defined in an analogous way, when the plane Π_α is the plane that contains v and is parallel to f, and f has to move inside Π_α. The case of a V-contact is analogous, and the plane Π_α is the plane containing the triangle of Γ that is involved in the contact.

When $\mathcal{P}$ is sliding along one contact α, it has three remaining degrees of freedom, two for moving along the plane Π_α, and one more for scaling. If we require $\mathcal{P}$ to maintain two contacts, α and β, we leave it with only two degrees of freedom. It is easy to see that the center of $\mathcal{P}$ is confined to move inside a specific plane in $\mathbb{R}^3$, and that this plane is incident to the line $\Pi_\alpha \cap \Pi_\beta$.

Sliding Π along three uni-contacts, α, β, and δ, means moving it in such way that it maintains all three contacts during the movement. This movement has only one degree of freedom, and the center of Π is confined to moving along a specific line $l_{\alpha,\beta,\delta}$ that is incident to the point $p_{\alpha,\beta,\delta} = \Pi_\alpha \cap \Pi_\beta \cap \Pi_\delta$ (in general position, this intersection is indeed a single point).

This means that if we wish to slide Π along three contacts as above, starting from a specific placement, we can do so in only two specific directions along $l_{\alpha,\beta,\delta}$. One of these directions brings $\mathcal{P}$ towards $p_{\alpha,\beta,\delta}$. Since $\mathcal{P}$ has to continue touching the planes Π_α, Π_β, and Π_δ during the movement, $\mathcal{P}$ shrinks to a point if the sliding continues until the center of $\mathcal{P}$ reaches $p_{\alpha,\beta,\delta}$. However, at least two of the contacts are necessarily lost prior to this point, since the triangles of Γ that define α, β, and δ are disjoint and only at most one of them can be incident to $p_{\alpha,\beta,\delta}$. This implies that at some point during the sliding towards $p_{\alpha,\beta,\delta}$, $\mathcal{P}$ will reach the boundary of one of the triangles it touches, and will lose the contact defined by this triangle completely if it continues sliding. At this moment, a bi-contact event will occur.

Because of linearity of trajectories and the preceding discussion, the volume of Π increases (resp., decreases) when $\mathcal{P}$ slides away from (resp., towards) $p_{\alpha,\beta,\delta}$. (In general, $\mathcal{P}$ does not shrink 'into itself' as it approaches $p_{\alpha,\beta,\delta}$, but sweeps new portions of space while shrinking.)

Consider the situation where we slide $\mathcal{P}$ along three contacts starting from an event X defined by four contacts. The general position assumption implies that the fourth contact made by $\mathcal{P}$ in the event will be lost immediately after the beginning of the sliding. Moreover, because of the linearity of trajectories, in one of the directions of the sliding this fourth contact will penetrate $\mathcal{P}$, while in the other direction, the placement of $\mathcal{P}$ will initially be free.

Suppose $\mathcal{P}$ makes the contacts α, β, and δ at some placement. Consider a face f of $\mathcal{P}$, such that one of the *open* half-spaces defined by the plane incident to f (at this placement) contains $p_{\alpha,\beta,\delta}$, but does not contain any point of $\mathcal{P}$. Define $\mathcal{P}^+ \subseteq \partial\mathcal{P}$ as the collection of such faces f, together with the edges and vertices of $\mathcal{P}$ incident to these faces. (These are the features of $\mathcal{P}$ that $p_{\alpha,\beta,\delta}$ 'sees' at this placement of $\mathcal{P}$.) Consider now a face g of $\mathcal{P}$, such that one of the *open* half-spaces defined by the plane incident to g contains neither $p_{\alpha,\beta,\delta}$ nor any point of $\mathcal{P}$. Define $\mathcal{P}^- \subseteq \partial\mathcal{P}$ as the collection of such faces g, together with the edges and vertices of $\mathcal{P}$ incident to these faces. (These are the features 'hidden' from $p_{\alpha,\beta,\delta}$.)

Observe that if $p_{\alpha,\beta,\delta}$ lies on a plane that is incident to a face h of $\mathcal{P}$, h belongs neither to $\mathcal{P}^+$ nor to $\mathcal{P}^-$. If $p_{\alpha,\beta,\delta}$ lies on a line that is incident to an edge e of $\mathcal{P}$, the same holds for e. If the point $p_{\alpha,\beta,\delta}$ is incident to a vertex v of $\mathcal{P}$, this also holds for v. Moreover, if $p_{\alpha,\beta,\delta}$ is incident to a vertex v of $\mathcal{P}$, then $\mathcal{P}^+ \equiv \emptyset$.

Observation.

- $\mathcal{P}^+$ and $\mathcal{P}^-$ are the same for any placement of $\mathcal{P}$ that makes the contacts α, β, and δ.
- If a new clean uni-contact is made by $\mathcal{P}$ while $\mathcal{P}$ slides towards $p_{\alpha,\beta,\delta}$ as above, such that the new contact involves a feature t of $\mathcal{P}$ and a triangle γ of Γ that was disjoint from $\mathcal{P}$ immediately before the contact was made, then t necessarily belongs to $\mathcal{P}^+$. If such situation occurs while $\mathcal{P}$ slides away from $p_{\alpha,\beta,\delta}$, then t belongs to $\mathcal{P}^-$.

We say that each face of $\mathcal{P}$ that belongs to $\mathcal{P}^+$ is a *frontier face*, and that $\mathcal{P}^+$ as a whole is the *frontier*, when we slide towards $p_{\alpha,\beta,\delta}$ as above. Symmetrically, each face of $\mathcal{P}$ that belongs to $\mathcal{P}^-$ is a frontier face and $\mathcal{P}^-$ is the frontier when we slide along α, β, and δ in the opposite direction.

Shrinking $\mathcal{P}$ into itself. Consider a free placement of $\mathcal{P}$ that makes three uni-contacts α, β, and δ, all incident to a vertex v of $\mathcal{P}$. It is easy to see that in this case $p_{\alpha,\beta,\delta} = v$, and consequently $\mathcal{P}^+ \equiv \emptyset$. (A 2-dimensional equivalent of this situation is illustrated in Figure 1(f).) This implies that if we slide $\mathcal{P}$ towards $p_{\alpha,\beta,\delta}$, along α, β, and δ, $\mathcal{P}$ will not encounter new uni-contacts and no uni-contact event will occur to $\mathcal{P}$. (Informally, $\mathcal{P}$ shrinks 'into itself'.) However, as observed above, $\mathcal{P}$ will at some point reach a bi-contact event. This implies that we can uniquely charge a free uni-contact event that has a vertex of degree 3 to a free bi-contact event.

2.4 Notation

We will distinguish between several types of events, based on the type of contacts that are involved in the event. Each type will be denoted mnemonically by listing the types of the involved contacts. For instance, an $FFEV$ event is an event with two F-contacts, an E-contact, and a V-contact. Note that it can equivalently be called an $EFVF$ event, say.

Every such type of events can have many combinatorially different sub-types. For instance, one $EEEE$ event may have 4 active vertices (the minimum possible number), while another may have 8 (the maximum possible number). Such combinatorially different sub-types of events will often get different treatment in our analysis, and it is thus essential to be able to distinguish between them in the text. However, for the sake of clarity in the exposition, we will not introduce a text-based notation to distinguish between them, but instead provide a symbolic illustration of each sub-type whose analysis is non-trivial (see, e.g., Figure 2), and refer to the relevant illustration when discussing a certain sub-type.

Each combinatorial type of events is represented for the purpose of illustration as a multi-hypergraph, whose vertices correspond to the vertices of $\mathcal{P}$ that are active in the event. This hypergraph may have edges that connect 1, 2, or 3 vertices, and such edges correspond respectively to V-, E- and F-contacts. The degree of a vertex in the hypergraph is thus the same as the degree of the corresponding vertex of $\mathcal{P}$ in an event of the illustrated combinatorial type. The illustrations show planar realizations of the above multi-hypergraphs, showing edges of degree 3 (F-contacts) as filled triangles, edges of degree 2 (E-contacts) as straight or curved segments, and edges of degree 1 (V-contacts) as small circles. For a concrete example, refer to Figure 3 that shows a specific combinatorial type of $FEVV$ events. In this type, the F-contact is incident to the E-contact, and the two V-contacts are incident to the F-contact and the E-contact, respectively.[3]

2.5 Illustration

Figure 1 provides an illustration in a 2-dimensional setting of some of the concepts introduced above. In the figure, the polygon $\mathcal{P}$ is shown with its center marked by a thick dot. Notice that in 2-D, a uni-contact event has only 3 contacts, and that $\mathcal{P}$ has one degree of freedom if it has to maintain 2 uni-contacts or one bi-contact. A free uni-contact event is shown in (a), while pseudo-free and 1-level free events are shown in (b) and (c), respectively. Although $\mathcal{P}$ has 6 vertices, only 4 are active in these events. In (d), $\mathcal{P}$ makes the contacts a and b, and can maintain these two contacts while sliding its center along the line $l_{a,b}$. The thickened edge is the only part of $\partial\mathcal{P}$ in $\mathcal{P}^+$, and if $\mathcal{P}$ slides towards the point $p_{a,b}$ and encounters a uni-contact event, the new third contact can only involve this edge. In (e), a bi-contact a is shown, together with the line l_a that $\mathcal{P}$ can slide its center along while maintaining this contact. In (f), $\mathcal{P}$ makes two contacts a and b that are incident to a vertex that has degree 2. In this case, $\mathcal{P}^+$ is empty, and $\mathcal{P}$ cannot encounter new uni-contact events when sliding towards $p_{a,b}$; the first event encountered during such sliding is necessarily a multi-contact event. (Informally, when sliding towards $p_{a,b}$, $\mathcal{P}$ shrinks 'into itself'.)

[3]Due to lack of space, we omit the treatment of triangle sites; hence the main figures do not depict V-contacts.

3. TECHNIQUES

In this section we outline the techniques and arguments we will be employing repeatedly throughout the analysis. Let Γ be a set of n pairwise disjoint triangles in $\mathbb{R}^3$, in general position.

3.1 Multi-contact events

The proofs of the following two lemmas are similar to parts of the analysis of Chew et al. [9], and are omitted due to space limitations.

LEMMA 3.1. *The number of free and pseudo-free events that involve a V^3-contact is $O(n)$.*

LEMMA 3.2. *The number of free and pseudo-free events that involve a V^2- or an E^2-contact is $O(n^2\alpha(n))$.*

We will only consider uni-contact events in the remainder of this section.

3.2 Popular Vertices

A vertex v of $\mathcal{P}$ is said to be *popular* in a certain event of $\mathcal{P}$ if its degree in this event is at least 3. A 2-dimensional equivalent of a popular vertex is shown in Figure 1(f). The proof of the following lemma is based on ideas introduced by Chew et al. [9].

LEMMA 3.3. *The number of free and pseudo-free events that are such that one of the vertices of $\mathcal{P}$ is popular is $O(n^2\alpha(n))$.*

PROOF. Consider a free event X, in which a vertex v of $\mathcal{P}$ is popular. Shrink $\mathcal{P}$ into itself towards v, and uniquely charge X to a free bi-contact event Y, as described at the end of Section 2.3. The lemma follows from the fact that the number of such events Y is $O(n^2\alpha(n))$, as stated in Lemma 3.2. Pseudo-free events are handled similarly. □

3.3 Popular Faces

A face f of $\mathcal{P}$ is said to be popular if each edge adjacent to it is either involved in an E-contact, or is incident to a face other than f that is involved in an F-contact. A 2-dimensional equivalent of a popular face is shown in Figure 1(d).

LEMMA 3.4. *The number of free events that are such that one of the faces of $\mathcal{P}$ is popular is $O(n^2\alpha(n))$.*

PROOF. Consider a free event X, in which a face f of $\mathcal{P}$ is popular. If one of the vertices incident to f is popular, then the lemma follows from Lemma 3.2. We thus assume below that the three vertices adjacent to f all have degree at most 2. The popularity of the face f implies that their degrees have to equal 2. This implies that exactly three contacts involve edges of f and faces adjacent to these edges. Let α, β, and δ be these three contacts. Consider sliding along α, β, and δ, as described in Section 2.3. It is easy to see that in one of the two possible directions of the sliding, f is the sole frontier face. Slide in this direction and consider the first event Y encountered during the sliding. We charge X to Y. As in the proof of Lemma 3.3, the charging is unique. Y is either a bi-contact event, or it is an event that involves α, β, δ, and a fourth contact that involves f or one of the edges or vertices of f. In the latter case, Y necessarily has a popular vertex. The lemma thus follows from the bounds stated in Lemmas 3.2 and 3.3. □

3.4 Reduction to Point Sites

LEMMA 3.5. *The number of free and pseudo-free $FFFF$ events of a polyhedron $\mathcal{P}$ with a constant number of vertices is $O(n^2)$.*

PROOF. Consider the set of vertices of the polygons in Γ. It is a set of $O(n)$ points in 3-space. It is easy to see that each free $FFFF$ event uniquely corresponds to a vertex in the Voronoi diagram of this set of points, under the distance function induced by $\mathcal{P}$. A pseudo-free $FFFF$ event similarly corresponds to a vertex in this diagram whose level is at most 8. The complexity of this diagram is $O(n^2)$ [14], and this is also a bound for the number of free and pseudo-free $FFFF$ events. □

3.5 Induction

Let $N_i = N_i(\Gamma)$ denote the maximum number of vertices of $Vor_{\mathcal{P}}(\Gamma)$, over all polytopes $\mathcal{P}$ with $\mathcal{P}$ has i vertices. It was observed in Section 2.2 (paragraph 'Activeness') that it suffices to consider i up to 12. The discussion in that section also implies that the number of events in which only $j < i$ vertices of $\mathcal{P}$ are active is $O(N_j)$, even if $\mathcal{P}$ has i vertices.

Our basic approach to deriving a bound on N_{12} is to use induction. We first prove a near-quadratic bound on N_4, for which it suffices to consider tetrahedral $\mathcal{P}$. We then bound N_i in terms of N_{i-1}, for $5 \le i \le 12$. One of the ways to achieve this will be to charge events with i active vertices to events with at most $i-1$ active vertices (see, e.g., Lemmas 3.6 and 3.8).

3.6 The SFC Technique, or Sliding away from a Face Contact

Consider a free uni-contact event X in which a face f of $\mathcal{P}$ is involved in an F-contact. Let α, β, and δ be the other three contacts involved in X. Consider sliding along α, β, and δ. By definition, only one of the parts $\mathcal{P}^+$ and $\mathcal{P}^-$ includes f. If this part is $\mathcal{P}^+$, we slide away from $p_{\alpha,\beta,\delta}$, and if this part is $\mathcal{P}^-$, we slide towards $p_{\alpha,\beta,\delta}$. It is easy to see that $\mathcal{P}$ is in a free placement immediately after the beginning of the sliding. We charge X to the first event Y encountered during the sliding. Y is necessarily free. Observe that during this sliding, the face f cannot become involved in new F-contacts, as observed in Section 2.3. Therefore, the fourth contact in Y does not involve the face f.

The SFC technique is useful when we want to charge an event that involves a particular F-contact to an event that does not involve an F-contact with the same face. Consider, for example, the following situation.

We say that an F-contact is 'isolated' if it is incident to no other contact. Suppose X involves an isolated F-contact, and denote the face of $\mathcal{P}$ involved in this contact by f. Use the SFC technique to slide away from this contact and charge X to an event Y. We claim that Y has fewer active vertices than X. Indeed, any contact that does not involve the face f cannot increase the degree of all the three vertices of f. Since the degree of these three vertices was 1 in X, at least one of these vertices has degree 0 in Y. This implies the following lemma.

LEMMA 3.6. *The number of free events of $\mathcal{P}$ that have an isolated F-contact is at most $O(N_{i-1})$, where i is the number of vertices in $\mathcal{P}$.*

3.7 The SEC Technique, or Sliding away from an Edge Contact

Consider a free event X in which a face f of $\mathcal{P}$ is involved in an F-contact α, and an edge e of f is involved in an E-contact. The degree of e is thus at least 2. Let β and δ be the other two contacts involved in X. The point $p_{\alpha,\beta,\delta}$ lies on the plane Π_α that is incident to the face f (at this placement). Thus, the edge e that lies inside this plane belongs to only one of the parts $\mathcal{P}^+$ and $\mathcal{P}^-$ (this holds for the other two edges of f as well). Sliding along α, β, and δ, as in the previous subsection, we can uniquely charge X to a free event Y that involves the contacts α, β, δ, and a fourth contact that is not an E-contact with the edge e. It is also easy to see that this fourth contact cannot be an F-contact with the second face incident to e. The ability to perform such charging will prove useful in several steps of our analysis.

3.8 Pseudo-free Events

Let i be the number of vertices in $\mathcal{P}$. Let M_i denote the maximum number of pseudo-free events of $\mathcal{P}$ among Γ, over all $\mathcal{P}$ with i vertices. The proof of the following lemma is omitted due to space limitations.

LEMMA 3.7. $M_i = O(N_i)$.

3.9 The LEM Technique, or Lower Envelopes Merging

Let i be the number of vertices in $\mathcal{P}$. Consider four (not necessarily distinct) features (each being a vertex, an edge, or a 2-dimensional face) of $\mathcal{P}$, such that two of these features, a, b, are incident to each other, and the other two, c, d, are each incident to a vertex of degree 1 (that is, no other feature is incident to this vertex). Consider a uni-contact event in which each of these features is involved in a distinct contact. Such events are said to be LEM events.

LEMMA 3.8. *The number of free and pseudo-free LEM events of $\mathcal{P}$ is $O(N_{i-1})$.*

The proof of the lemma is based on the proof of [9, Lemma 3.5]. Informally, fixing the contacts a, b leaves $\mathcal{P}$ with 2 degrees of freedom, so its corresponding placements can be represented within a 2-dimensional frame, in which the events under consideration can be shown to be vertices of the lower envelope of all the c-contacts and the d-contacts. We then use the fact that the complexity of such an envelope is proportional to the sum of the complexities of the two sub-envelopes of the c-contacts and of the d-contacts.

3.10 The Tagansky Technique

Consider four (not necessarily distinct) features a, b, c, and d (each being a vertex, an edge, or a 2-dimensional face) of $\mathcal{P}$. Consider all the uni-contact events in which each of these features is involved in a distinct contact. Denote the collection of such free events by $\Omega_{a,b,c,d}$, and the collection of such 1-level events by $\Omega^1_{a,b,c,d}$. Let Φ be another collection of events, such that $|\Phi| = O(n^2\alpha(n)\log^c n)$, for some non-negative integer constant c.

Assume that, given an event of $\Omega_{a,b,c,d}$, we can either **(i)** charge it to $i > 0$ events of $\Omega^1_{a,b,c,d}$, such that the number of times that each event of $\Omega^1_{a,b,c,d}$ is charged in this fashion is at most j, or **(ii)** charge it to an event of Φ, such that the number of times that each event of Φ is charged in this fashion is bounded by a constant. The charging is typically done by sliding on three of the contacts and letting the fourth contact penetrate $\mathcal{P}$. The following lemma is directly implied by the work of Tagansky [20], and its proof is omitted due to space limitations.

LEMMA 3.9. *Under the assumptions described above:*

- *If $i/j = 2$, then $|\Omega_{a,b,c,d}| = O(n^2\alpha(n)\log^{c+1} n)$.*
- *If $i/j > 2$, then $|\Omega_{a,b,c,d}| = O(n^2\alpha(n)\log^{c} n)$.*

4. SEGMENT SITES

THEOREM 4.1. *The complexity of the Voronoi diagram of a set Γ of n segments in 3-space, under a convex distance function induced by a polytope $\mathcal{P}$ with q facets, is $O(q^4 n^2\alpha(n)\log n)$, provided that the segments are in general position with respect to $\mathcal{P}$, as defined in Section 2.1.*

PROOF. Let i be the number of vertices of $\mathcal{P}$. As described in Section 3.5, the proof of Theorem 4.1 proceeds by induction on i, through a series of lemmas. Recall that it suffices to confine ourselves to $i \leq 12$. We begin by obtaining a bound for N_2 (Lemma 4.2), and then work our way up to N_{12} (Lemma 4.9). Notice that the only uni-contacts that occur are F- and E-contacts. Recall that it is sufficient to bound the number of free events.

A free or pseudo-free event (or a 1-level free or pseudo-free event) is said to be *good* if it is a multi-contact event or an $FFFF$ event, or if it contains a popular vertex or an isolated F-contact. A free event (or a 1-level free event) that contains a popular face is also said to be good. Events with isolated F-contacts are easily handled using Lemma 3.6, and the various results shown in Section 3 imply that the number of free and pseudo-free good events without isolated F-contacts is $O(n^2\alpha(n))$. We thus do not treat good events explicitly below.

LEMMA 4.2. *The complexity of the Voronoi diagram of a set of segments in 3-space is $O(n^2)$ under a distance function induced by a segment, and is $O(n^2\alpha(n))$ under a distance function induced by a triangle*[4]*. In other words, $N_2 = O(n^2)$ and $N_3 = O(n^2\alpha(n))$.*

PROOF. It is easy to see that if $\mathcal{P}$ is a triangle (3 vertices), all the events of $\mathcal{P}$ among Γ either have a popular vertex or a multi-contact. A bound of $O(n^2\alpha(n))$ on the complexity of $Vor_{\mathcal{P}}(\Gamma)$ in these cases is thus implied by Lemmas 3.3, 3.1 and 3.2. If $\mathcal{P}$ is a segment (2 vertices), it is easy to see that, assuming general position, the only type of events is V^2V^2, and each event can be uniquely charged to a pair of segments of Γ. This easily implies a bound of $O(n^2)$. ☐

LEMMA 4.3. *The complexity of the Voronoi diagram of a set of segments in 3-space is $O(n^2\alpha(n)\log n)$ under a distance function induced by a tetrahedron. In other words, $N_4 = O(n^2\alpha(n)\log n)$.*

[4]Strictly speaking, these are not well-defined distance functions. What the lemma actually analyzes is the number of free events of the underlying segment or triangle, and it should be interpreted only in this context.

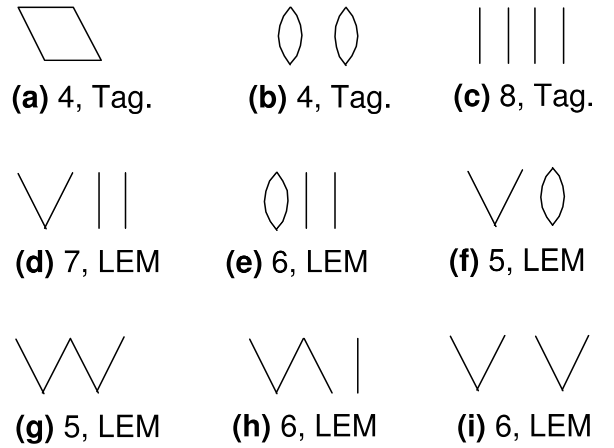

Figure 2: Different combinatorial types of $EEEE$ events, not including good events. For each type, the number of active vertices and the name of the technique employed in the analysis are stated.

PROOF. Let $\mathcal{P}$ be a tetrahedron. Notice that all events of $\mathcal{P}$ that have at least one F-contact, have at least one popular vertex. Indeed, observe that an E-contact raises the degree of two vertices of $\mathcal{P}$ by 1, while an F-contact similarly raises the degree of three vertices. Thus, the sum of the degrees of the four vertices of $\mathcal{P}$ in an event with one F-contact and three other contacts that are either E or F is at least $3 + 2 + 2 + 2 = 9$. However, this necessarily means that one of the four vertices of $\mathcal{P}$ has degree at least 3, and is therefore popular. Thus, the $O(n^2\alpha(n))$ bound on the number of events that have at least one F-contact follows from Lemma 3.3. This bound also applies to the number of $EEEE$ events that have a popular vertex.

We now bound the number of $EEEE$ events that are not good, and therefore have no popular vertices or faces. Figure 2 illustrates all the combinatorially distinct types of such events when $\mathcal{P}$ is a polyhedron with an arbitrary number of vertices. Only two of them, the ones shown in Figures 2(a) and 2(b), can occur when $\mathcal{P}$ is a tetrahedron, since all the other types have at least 5 active vertices. We treat these two types of events using the Tagansky technique (see Section 3.10). Due to lack of space, we merely state that the analysis shows that the number of events of the Figure 2(a) and Figure 2(b) types is $O(n^2\alpha(n)\log n)$. (A similar analysis was carried out by Chew et al. [9], but only for the case of line sites.) ☐

LEMMA 4.4. $N_5 = O(n^2\alpha(n)\log n)$.

PROOF. Let $\mathcal{P}$ be a convex polyhedron with 5 vertices. Events of $\mathcal{P}$ with 5 active vertices are either good, or belong to the types illustrated in Figures 2(f), 2(g), 4(m), 4(n), 4(o), 5(y) and 5(z). Events of the types shown in Figures 2(f) and 2(g) are LEM events, and their number is thus $O(N_4) = O(n^2\alpha(n)\log n)$ (Lemmas 3.8 and 4.3).

Let X be a free event of a type illustrated in Figure 4(n) or 4(o). Let f be the face of $\mathcal{P}$ that is involved in the F-contact in X. Use the SFC technique (Section 3.6) to slide away from this F-contact. Charge X to the first event Y encountered during the sliding. Y is a free event, since the placement of $\mathcal{P}$ immediately upon the beginning of the sliding is free. The SFC technique implies that either Y is a multi-contact event, or it has the same triple of E-contacts as X, together with a fourth contact that is not an F-contact with f. It is easy to see that either **(i)** Y has 4 active vertices, or **(ii)** Y is a good event, or **(iii)** Y is a LEM event. Since the number of such events Y is $O(n^2\alpha(n)\log n)$, and any event Y can be charged as above only a constant number of times, this implies that the number of events of the types illustrated in Figures 4(n) and 4(o) is $O(n^2\alpha(n)\log n)$.

Let X be a free event of the type illustrated in Figure 4(m). Let a denote the edge of $\mathcal{P}$ that is involved in the E-contact that is incident with two vertices to the F-contact. Use the SEC technique (Section 3.7) to slide away from this E-contact and charge a free event Y, as above. The SEC technique implies that Y cannot have an E-contact that involves a, or a second F-contact that involves a face incident to a. This implies that either **(i)** Y is a good event, or **(ii)** Y is an event of the type shown in Figure 4(o). The number of events of the Figure 4(m) type is therefore $O(n^2\alpha(n)\log n)$, as above.

Let X be a free event of one of the type illustrated in Figures 5(y) and 5(z). Use the SFC technique (Section 3.6) to slide away from one of the two F-contacts in X. Charge X to the first event Y encountered during the sliding. As above, Y is a free event, and does not involve an F-contact with the same face. Thus, either **(i)** Y is a good event, or **(ii)** Y is an event of one of the types shown in Figures 4(m), 4(n) and 4(o). The number of events of the types illustrated in Figures 5(y) and 5(z) is thus $O(n^2\alpha(n)\log n)$. □

LEMMA 4.5. *$N_6 = O(n^2\alpha(n)\log n)$.*

PROOF. Let $\mathcal{P}$ be a convex polyhedron with 6 vertices. Events of $\mathcal{P}$ with 6 active vertices are either good, or belong to the types illustrated in Figures 2(e), 2(h), 2(i), 4(b), 4(d), 4(f), 4(g), 4(i), 4(j), 4(k), 5(g), 5(h), 5(l), 5(s), 5(u), 5(v), 5(A), 5(B), 5(C), 6(q), and 6(r).

Free events of the types shown in Figures 2(e), 2(h), 2(i), 4(b), 4(d), 4(f), 4(g), 4(i), 4(j), 4(k), 5(g), 5(h), 5(l), 5(s), 5(u), and 5(v) are LEM events, and their number is thus $O(N_5) = O(n^2\alpha(n)\log n)$ (Lemmas 3.8 and 4.4).

Let X be a free event of a type illustrated in Figure 5(A), 5(B), or 5(C). Use the SFC technique to slide away from any one of the two F-contacts of X and charge X to a free event Y, as in the proof of Lemma 4.4. Either **(i)** Y is good, or **(ii)** Y is a LEM event, or **(iii)** Y has at most 5 active vertices. The number of events of the types illustrated in Figures 5(A), 5(B), and 5(C) is therefore $O(n^2\alpha(n)\log n)$.

Let X be a free event of a type illustrated in Figure 6(q) or 6(r). Use the SFC technique to slide away from the (unique) F-contact that involves a face that is incident to a vertex of degree 1. Charge X to the first event Y encountered during the sliding. In complete analogy to the above, the number of these events is $O(n^2\alpha(n)\log n)$. This completes the proof of the lemma. □

LEMMA 4.6. *$N_7 = O(n^2\alpha(n)\log n)$.*

PROOF. Let $\mathcal{P}$ be a convex polyhedron with 7 vertices. Events of $\mathcal{P}$ with 7 active vertices are either good, or belong to the types illustrated in Figures 2(d), 4(a), 4(e), 4(h), 4(l), 5(a), 5(c), 5(d), 5(e), 5(j), 5(k), 5(m), 5(n), 5(p), 5(r), 5(t), 5(x), 6(c), 6(h), 6(j), 6(k), 6(l), 6(n), and 6(o). All of these events are LEM events, and the number of free such events is thus $O(N_6) = O(n^2\alpha(n)\log n)$ (Lemmas 3.8 and 4.5). □

LEMMA 4.7. *$N_8 = O(n^2\alpha(n)\log n)$.*

PROOF. Let $\mathcal{P}$ be a convex polyhedron with 8 vertices. Events of $\mathcal{P}$ with 8 active vertices are either good, or belong to the types illustrated in Figures 2(c), 4(c), 5(b), 5(f), 5(o), 5(q), 5(w), 6(a), 6(b), 6(e), 6(g), 6(i), 6(m), and 6(p). All of these events except the Figure 2(c) type are LEM events, and the number of free such events is $O(N_7) = O(n^2\alpha(n)\log n)$ (Lemmas 3.8 and 4.6).

We treat events of the type illustrated in Figure 2(c) using the Tagansky technique (see Section 3.10). Due to lack of space, we omit all details of this analysis, that shows that the number of events of the Figure 2(c) type is $O(n^2\alpha(n)\log n)$. □

LEMMA 4.8. *$N_9 = O(n^2\alpha(n)\log n)$.*

PROOF. Let $\mathcal{P}$ be a convex polyhedron with 9 vertices. Events of $\mathcal{P}$ with 9 active vertices are either good, or belong to the types illustrated in Figures 5(i), 6(d), and 6(f). All of these events are LEM events, and the number of free such events is $O(N_8) = O(n^2\alpha(n)\log n)$ (Lemmas 3.8 and 4.5). □

LEMMA 4.9. *$N_d = O(n^2\alpha(n)\log n)$, for $10 \le d \le 12$.*

PROOF. A free event of $\mathcal{P}$ with 10 or more active vertices necessarily has at least one isolated F-contact. The number of such events is therefore easily bounded by induction, using Lemma 3.6. Since $N_9 = O(n^2\alpha(n)\log n)$ (Lemma 4.8), this implies the lemma. □

This completes the proof of Theorem 4.1. □

5. POLYHEDRAL SITES

THEOREM 5.1. *The complexity of the Voronoi diagram of a collection of pairwise disjoint polyhedral sites in 3-space that have n vertices overall, under a convex distance function induced by a polytope $\mathcal{P}$ with q facets, is $O(q^4 n^{2+\varepsilon})$, for any $\varepsilon > 0$, provided that the sites are in general position with respect to $\mathcal{P}$, as defined in Section 2.1.*

Unfortunately, all details of the proof of Theorem 5.1 have to be omitted in this version due to space limitations. The proof proceeds by induction on the number of vertices of $\mathcal{P}$, as in Section 4, starting with $\mathcal{P}$ being a segment or a triangle, and concluding with the case of $\mathcal{P}$ having 12 vertices.

The proof is complicated by the fact that V-contacts, which cannot occur when the sites are segments, can now appear. This raises the number of combinatorially distinct types of events that have to be handled, and drastically increases the number of such types that cannot be handled 'easily' with the LEM, the SFC, and the SEC techniques.

The hardest stages of the proof are bounding N_4 and N_5. For N_4, we resort to the analysis technique of counting schemes, introduced by Halperin and Sharir [12, 18], and refined in several subsequent papers (see, e.g., [1, 15]). Informally, in the refined version, we charge each 0-level event to about k^2 events at level at most k (i.e., at most k sites intersect the interior of $\mathcal{P}$ at such events), or to other events whose number can be bounded independently. This leads to a recurrence whose solution is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. (Actually, only one combinatorial type, shown in Figure 3, requires, so far, the use of a counting scheme.)

For N_5, we apply an intricate geometric analysis that strongly relies on the properties of the pentahedron—the only combinatorial form that a convex polytope with 5 vertices can assume. The Tagansky technique is employed repeatedly, sometimes in a fairly involved fashion, throughout the proof.

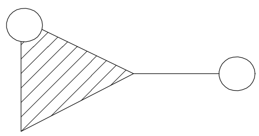

Figure 3: The 'hard' combinatorial type of events whose analysis involves a counting scheme.

6. APPROXIMATE NEAREST-NEIGHBOR SEARCHING

Theorem 5.1 can be applied to obtain the following result.

THEOREM 6.1. *We can preprocess a collection of disjoint polyhedra in 3-space with n vertices altogether into a data structure of size $O(n^{2+\varepsilon}/\delta^4)$, for any $\varepsilon > 0$, such that this data structure can answer δ-approximate Euclidean nearest-neighbor queries amidst the polyhedra in time $O(\log(n/\delta))$, for an arbitrarily small $\delta > 0$.*

The data structure described in the theorem is essentially a point-location data structure on a polyhedral Voronoi diagram of the collection of polyhedra ('sites'). We use the fact that the Euclidean ball in $\mathbb{R}^3$ can be δ-approximated by a convex polytope with $O(1/\delta)$ vertices. The polyhedral Voronoi diagram of the sites under the distance function induced by this polytope is a δ-approximation of the Euclidean Voronoi diagram of these sites. δ-approximate Euclidean nearest-neighbor queries amidst the sites can therefore be answered using point location queries in this polyhedral Voronoi diagram. Theorem 5.1 states that the complexity of this diagram is $O(n^{2+\varepsilon}/\delta^4)$, and standard machinery can be used to preprocess it into a point location data structure with the desired performance. All further details are omitted.

REMARK 1. *We also have another solution, with a data structure of size only $O(n^{2+\varepsilon}/\delta)$. However, the query time becomes $O((\log n)/\delta)$.*

Acknowledgements

Part of the work on this paper has been carried out while the first author was visiting the University of Illinois at Urbana-Champaign, hosted by Sariel Har-Peled. The author is grateful to Sariel for the hospitality and for many helpful discussions. In particular, the results reported in Section 6 were developed in collaboration with Sariel.

7. REFERENCES

[1] P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete Comput. Geom.*, 15:1–13, 1996.

[2] P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[3] P. K. Agarwal and M. Sharir. Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. *Discrete Comput. Geom.*, 24:645–685, 2000.

[4] B. Aronov and M. Sharir. On translational motion planning of a convex polyhedron in 3-space. *SIAM J. Comput.*, 26:1785–1803, 1997.

[5] B. Aronov, M. Sharir, and B. Tagansky. The union of convex polyhedra in three dimensions. *SIAM J. Comput.*, 26:1670–1688, 1997.

[6] 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.

[7] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete Comput. Geom.*, 19(4):473–484, 1998.

[8] L. P. Chew. Near-quadratic bounds for the L_1 Voronoi diagram of moving points. *Comput. Geom. Theory Appl.*, 7:73–80, 1997.

[9] L. P. Chew, K. Kedem, M. Sharir, B. Tagansky, and E. Welzl. Voronoi diagrams of lines in 3-space under polyhedral convex distance functions. *J. Algorithms*, 29(2):238–255, 1998.

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

[11] 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.

[12] D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions, with applications to visibility in terrains. *Discrete Comput. Geom.*, 12:313–326, 1994.

[13] S. Har-Peled. A replacement for Voronoi diagrams of near linear size. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, pages 94–103, 2001.

[14] 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.

[15] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. In *Proc. 42nd IEEE Symposium on Foundations of Computer Science*, pages 56–65, 2001.

[16] V. Koltun and M. Sharir. Three dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations. In *Proc. 18th ACM Symposium on Computational Geometry*, 2002.

[17] M. I. Shamos and D. Hoey. Closest-point problems. In *Proc. 16th IEEE Symposium on Foundations of Computer Science*, pages 151–162, 1975.

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

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

[20] B. Tagansky. *The Complexity of Substructures in Arrangements of Surfaces.* Ph.D. thesis, Tel Aviv University, Tel Aviv, 1996.

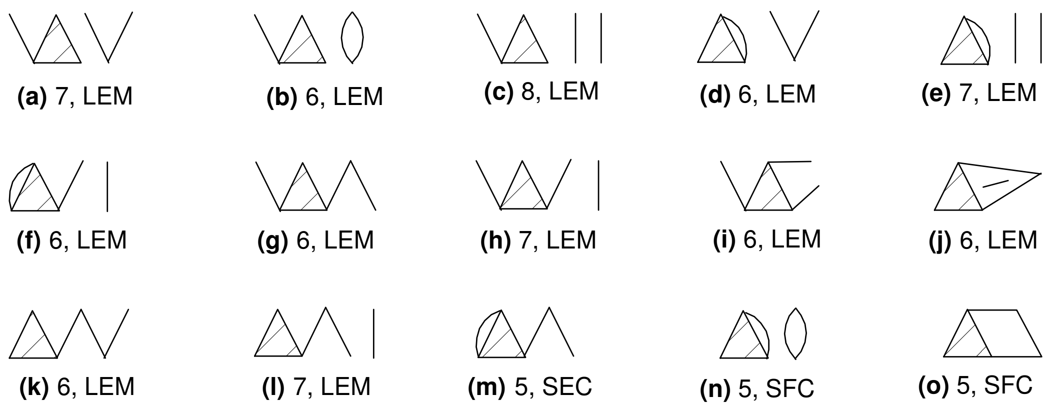

Figure 4: Different combinatorial types of $FEEE$ events, not including good events. For each type, the number of active vertices and the name of the technique employed in the analysis are stated.

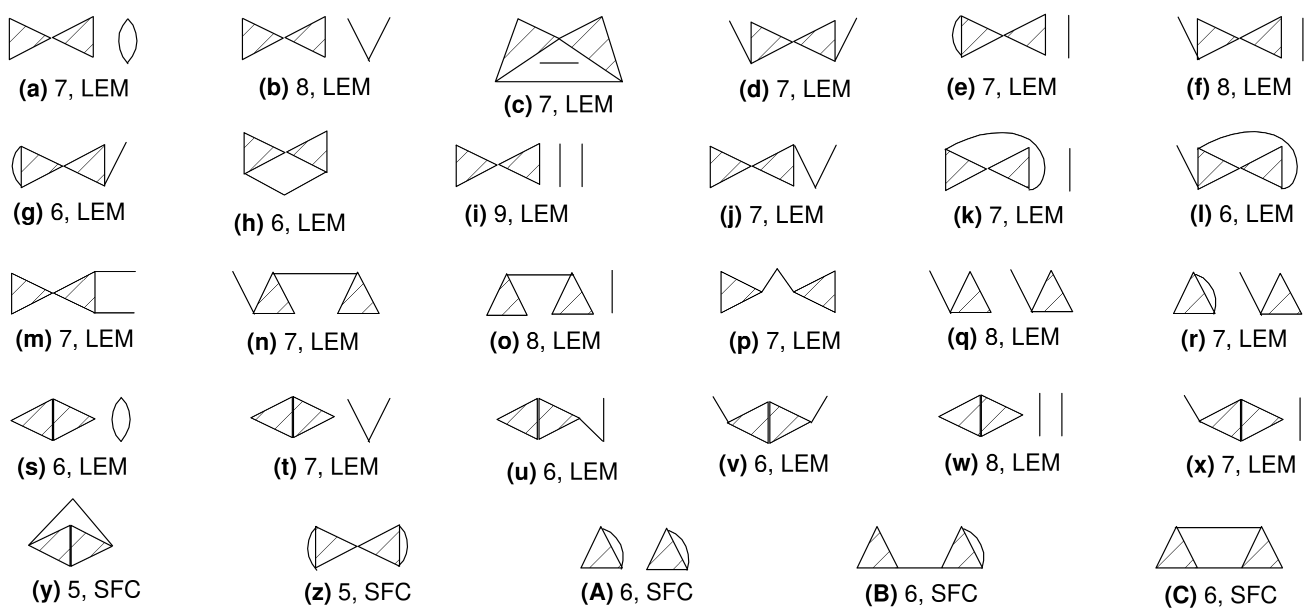

Figure 5: $FFEE$ events.

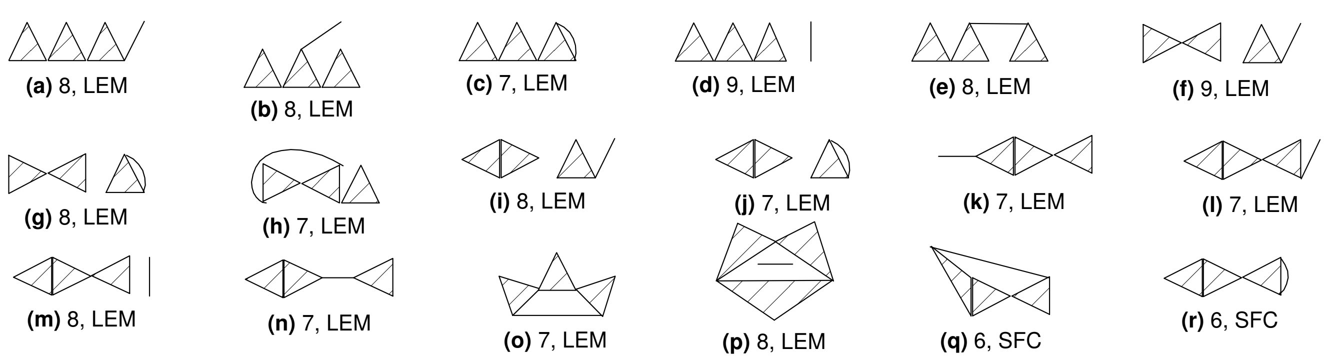

Figure 6: $FFFE$ events.