

Sharp Bounds for Vertical Decompositions of Linear Arrangements in Four Dimensions*

Vladlen Koltun

Computer Science Division, University of California,
Berkeley, CA 94720-1776, USA
vladlen@cs.berkeley.edu

Abstract. We prove tight and near-tight combinatorial complexity bounds for vertical decompositions of arrangements of hyperplanes and 3-simplices in four dimensions. In particular, we prove a tight upper bound of $\Theta(n^4)$ for the vertical decomposition of an arrangement of n hyperplanes in four dimensions, improving the best previously known bound [8] by a logarithmic factor. We also show that the complexity of the vertical decomposition of an arrangement of n 3-simplices in four dimensions is $O(n^4\alpha(n)\log^2 n)$, where $\alpha(n)$ is the inverse Ackermann function, improving the best previously known bound [2] by a near-linear factor.

1. Introduction

Given a collection Γ of n fixed-degree algebraic surfaces in \mathbb{R}^d , its arrangement [1], [9] is denoted by $\mathcal{A}(\Gamma)$. The complexity of a single cell in $\mathcal{A}(\Gamma)$ can easily reach $\Omega(n^{d-1})$, and may even be slightly super- $O(n^{d-1})$, depending on the degree of the surfaces of Γ [1]. There exist many geometric algorithms that require arrangements to be decomposed into cells of constant description complexity (that is, defined in terms of a constant number of polynomial equations and inequalities of constant maximum degree). As a result, devising a decomposition scheme that decomposes arrangements into as few as possible such cells is an intensively studied problem.

The most efficient general-purpose decomposition scheme is *vertical decomposition*, which is suitable for arrangements of fixed-degree algebraic surfaces in any dimension.

* This work was carried out while the author was a Ph.D. student in Tel Aviv University and was supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing). A preliminary version of this paper appeared in the *Proceedings of the 7th Workshop on Algorithms and Data Structures*, 2001, pp. 99–110.

We define four-dimensional vertical decompositions in Section 2, and refer the reader to [2] for details on the general scheme. Vertical decomposition was originally introduced in the context of two-dimensional problems, and was extended to higher dimensions in the late 1980s [2]. The complexity of the vertical decomposition of an arrangement of n triangles (or planes) in \mathbb{R}^3 is known to be $O(n^2\alpha(n)\log n + K)$, where K is the complexity of the undecomposed arrangement (which, in the worst case, is $O(n^3)$) [13]. However, there is still a substantial gap between the known upper and lower bounds for the complexity of vertical decompositions in dimensions higher than three.

The number of cells in the vertical decomposition of an arrangement of n fixed-degree algebraic surfaces in \mathbb{R}^4 has recently been shown, in a companion paper by the author [10], to be $O(n^{4+\varepsilon})$, for any $\varepsilon > 0$ (where the constant of proportionality depends on ε). This has improved a previous bound of $O(n^5\beta(n))$ [2], where $\beta(n)$ is an extremely slow-growing function of n (it is an exponential function of $\alpha(n)$, which also depends on d and on the maximum degree of the polynomials that define the surfaces of Γ), related to Davenport–Schinzel sequences [12]. The importance of the problem of improving this previous upper bound has been emphasized numerous times [1], [2], [8], [9], [12], but the problem remained open for more than a decade. The analysis in [10] modifies and extends the approach introduced and developed in the current work.

An interesting previously known result concerning vertical decompositions in four dimensions is due to Guibas et al. [8], who showed that the vertical decomposition of an arrangement of n hyperplanes in four dimensions has complexity $O(n^4 \log n)$.

In this paper we prove a tight bound of $\Theta(n^4)$ for the case of hyperplanes, just mentioned, thereby improving the result of Guibas et al. [8]. Moreover, we bound the complexity of the vertical decomposition of an arrangement of n 3-simplices in four dimensions by $O(n^4\alpha(n)\log^2 n)$. This improves the best previously known upper bound for this setting [2] by a near-linear factor.

We note that an arrangement of simplices as above can be decomposed into $\Theta(n^4)$ cells by extending the simplices into hyperplanes, and decomposing the resulting hyperplane arrangement using bottom-vertex simplicial decomposition [3]. In other words, for both cases of hyperplanes and simplices, alternative $\Theta(n^4)$ decompositions are available. Nevertheless, since vertical decomposition is the most efficient known general-purpose decomposition technique for general surfaces, it is interesting to study its behavior also for the simpler case of linear arrangements. Moreover, this study is important since it introduces general techniques for analyzing vertical decompositions in four dimensions.

As mentioned above, this work has indeed been successful in this respect, since the techniques introduced herein were subsequently modified and extended to produce the desired near-optimal bound for the case of fixed-degree algebraic surfaces [10]. Moreover, the results obtained in the current paper are also interesting in themselves, as they are considerably sharper than the $O(n^{4+\varepsilon})$ bound of [10]. In particular, the presented bound for hyperplanes is tight.

Furthermore, a variety of novel techniques applied in this paper were specifically tailored to produce sharp bounds for the case of linear surfaces, and are not immediately applicable in the case of curved algebraic surfaces. These techniques appear only in the current paper, while in the sequel [10] they are replaced by “messier” albeit more general analysis techniques.

2. Vertical Decompositions in Four Dimensions

2.1. The Construction

Denote the coordinates by x , y , z , and w ; we treat z as the “vertical” coordinate in \mathbb{R}^4 . Given a collection Γ of 3-simplices (or hyperplanes) in \mathbb{R}^4 , the vertical decomposition of $\mathcal{A}(\Gamma)$, denoted by $\mathcal{V}(\Gamma)$, is constructed as follows.

For each $S \in \Gamma$, erect a three-dimensional z -vertical visibility wall on the boundary of S (a two-dimensional piecewise linear surface denoted by ∂S), which is defined as the union of all z -vertical segments that have an end-point on ∂S and are interior-disjoint from all simplices of Γ ; we refer to a segment whose interior is disjoint from all simplices of Γ as a *visibility* segment. Also, for each pair $S, T \in \Gamma$, erect a three-dimensional z -vertical visibility wall on the two-dimensional (linear) surface $S \cap T$ in a similar fashion. This results in a decomposition of $\mathcal{A}(\Gamma)$ into (not necessarily convex) z -vertical prisms, such that the floor (respectively, ceiling) of each prism, if it exists, is contained in a single simplex of Γ . We denote this decomposition by $\mathcal{V}_1(\Gamma)$.

For a prism P of $\mathcal{V}_1(\Gamma)$, the simplex containing its floor (respectively, ceiling) is denoted by P_F (respectively, P_C). Projecting P onto its floor or ceiling (in the z -direction) results in a three-dimensional polyhedron, denoted by P_{3D} . We first decompose P_{3D} by erecting two-dimensional y -vertical visibility walls on each of its edges. For an edge f of P_{3D} , the wall erected on it is defined as the union of all y -vertical (visibility) segments that have an end-point on f and are fully contained in P_{3D} . We then decompose P by erecting z -vertical three-dimensional visibility walls on each such y -vertical two-dimensional visibility wall of P_{3D} . Repeating this process for all prisms P results in a decomposition of $\mathcal{V}_1(\Gamma)$, which we denote by $\mathcal{V}_2(\Gamma)$.

We further refine $\mathcal{V}_2(\Gamma)$ as follows. For a prism Q of $\mathcal{V}_2(\Gamma)$, consider its z -projection Q_{3D} (which is a y -vertical prism in \mathbb{R}^3). Projecting Q_{3D} onto its floor or ceiling (in the y -direction) results in a two-dimensional polygon Q_{2D} , which we decompose by erecting zero, one, or two x -vertical (possibly infinite) visibility segments (defined in analogy to the above) on each vertex v of Q_{2D} . We then erect y -vertical two-dimensional walls (inside Q_{3D}) on each such x -vertical segment of Q_{2D} . Subsequently, z -vertical three-dimensional walls (inside Q) are erected on each such y -vertical two-dimensional wall of Q_{3D} . Repeating this process for each prism Q of $\mathcal{V}_2(\Gamma)$ decomposes $\mathcal{V}_2(\Gamma)$ into cells of constant description complexity: each such cell is a convex polyhedron bounded by up to eight facets.

This completes the construction of the vertical decomposition $\mathcal{V}(\Gamma)$. If Γ consists of hyperplanes, the construction is simpler, because there are no boundary features present, and all cells that arise, in all stages of the construction, are convex polyhedra.

Similar constructions can be formulated for arrangements of general algebraic surfaces in \mathbb{R}^4 , and for arrangements of simplices or algebraic surfaces in higher dimensions. We refer the reader to [2] for more details on the vertical decomposition scheme in these more general settings.¹

¹ Note our nonstandard order of the coordinates: the last coordinate w is not involved in the construction, unlike in the approach of Chazelle et al. [2].

2.2. Preliminary Analysis

The following two observations will be helpful in analyzing the complexity of $\mathcal{V}(\Gamma)$. First, it is easy to see that the complexity of $\mathcal{V}(\Gamma)$ is asymptotically the same as the complexity of $\mathcal{V}_2(\Gamma)$, and it is thus sufficient to bound the latter in order to bound the former. Indeed, the last stage of the above construction involves vertical decompositions of two-dimensional polygons, which only increase the number of features by a constant factor.

Second, the complexity of $\mathcal{V}_2(\Gamma)$ is asymptotically the same as the complexity of $\mathcal{V}_1(\Gamma)$ plus the number of *y-vertical visibility events* inside all the projection polyhedra P_{3D} of the prisms P of $\mathcal{V}_1(\Gamma)$. Such an event is said to happen between two edges, f and f' , of P_{3D} if they intersect a common *y-vertical line* l , and the segment $s \subset l$ that connects f and f' on this line lies completely inside P_{3D} .

We start by bounding the complexity of $\mathcal{V}_1(\Gamma)$.

Lemma 2.1. *Given a collection Γ of n 3-simplices in \mathbb{R}^4 , the complexity of $\mathcal{V}_1(\Gamma)$ is $O(n^4\alpha(n))$. If Γ consists of hyperplanes, the complexity of $\mathcal{V}_1(\Gamma)$ is $O(n^4)$.*

Proof. During the construction of $\mathcal{V}_1(\Gamma)$, a *z-vertical visibility wall* is erected on $S \cap T$, for all $S, T \in \Gamma$. By construction, this wall is bounded from above by the lower envelope of the part of $\mathcal{A}(\Gamma)$ that lies above $S \cap T$ (within the *z-vertical hyperplane* spanned by $S \cap T$). Similarly, it is bounded from below by the upper envelope of the part of $\mathcal{A}(\Gamma)$ that lies below $S \cap T$. One such envelope can be alternatively viewed as follows: The cross section of a 3-simplex $R \in \Gamma$ with the *z-vertical hyperplane* spanned by $S \cap T$ is a triangle or a quadrilateral; its portion above (or below) the intersection $S \cap T$ is a convex k -gon, for $k = O(1)$; we are interested in the lower (or upper) envelope of the collection of these k -gons that correspond to all $R \in \Gamma$.

The complexity of each of these upper and lower envelopes is $O(n^2\alpha(n))$ [6]. Consequently, the complexity of the wall erected on $S \cap T$ is also $O(n^2\alpha(n))$. The same arguments imply that the complexity of the visibility wall erected on each ∂S (for $S \in \Gamma$) is $O(n^2\alpha(n))$ as well. Since there are n such boundaries ∂S and $O(n^2)$ such intersections $S \cap T$, the overall number of features that are created during the construction of $\mathcal{V}_1(\Gamma)$ is $O(n^4\alpha(n))$.

If Γ is a collection of hyperplanes, $S \cap T$ is a two-dimensional plane (for all $S, T \in \Gamma$), and the complexity of the above-mentioned envelopes is $O(n^2)$. This is because the envelopes are a portion of the zone of $S \cap T$ within the cross section of $\mathcal{A}(\Gamma)$ in the *z-vertical hyperplane* spanned by $S \cap T$, and the complexity of such a zone is $O(n^2)$ [7]. The number of created features in this case is thus $O(n^4)$. \square

2.3. Visibility Events

In light of the above, the bulk of this paper is devoted to bounding the number of *y-vertical visibility events* inside all polyhedra P_{3D} . To this end, we first classify the faces and edges of a projected prism P_{3D} , obtained as the *z-vertical projection* of a prism P of $\mathcal{V}_1(\Gamma)$. Each face of P_{3D} is a *z-vertical projection* of a face of P , which is a part of a

three-dimensional z -vertical wall erected on a two-dimensional feature of $\mathcal{A}(\Gamma)$ (during the construction of $\mathcal{V}_1(\Gamma)$). These walls can be of three types:

- An (upward) visibility wall erected on $P_F \cap S$ (for some $S \in \Gamma$) that touches P_C (from below). Faces of P_{3D} that are projections of parts of such walls are said to be *red*.
- A (downward) visibility wall erected on $P_C \cap T$ (for some $T \in \Gamma$) that touches P_F (from above). Corresponding faces of P_{3D} are said to be *green*.
- A visibility wall erected on ∂U (for some $U \in \Gamma$) that touches P_F (from above) and P_C (from below). (Intuitively, the boundary of U is partly “floating” between P_F and P_C and the z -vertical wall erected on it reaches both P_F and P_C .) Corresponding faces of P_{3D} are said to be *blue*. (Note that in the case of hyperplanes there are no blue faces.)

Any edge of P_{3D} is incident to two faces of P_{3D} . Such edges can thus be classified into four types, depending on the types of the incident faces:

1. Edges incident to two red (or two green) faces. The two faces, by definition, correspond to parts of the visibility walls erected on $P_F \cap S$ and $P_F \cap T$ (respectively, on $P_C \cap S$ and $P_C \cap T$), for some $S, T \in \Gamma$. An edge incident to both of them thus corresponds to the common part of these two walls, which is the visibility wall erected on $P_F \cap S \cap T$ (respectively, $P_C \cap S \cap T$). We denote such edges mnemonically by E_3 to signify that they are formed by an intersection of three simplices.
2. Edges incident to a red and a green face. Such an edge corresponds to the common part of two walls, one erected (upward) on $P_F \cap S$ and another erected (downward) on $P_C \cap T$, for some $S, T \in \Gamma$. This part is composed of z -vertical segments that touch both $P_F \cap S$ and $P_C \cap T$. We denote such edges by E_{22} to signify that they are formed by the interaction of two intersections, each of two simplices.
3. Edges incident to a red (or green) face and a blue face. Such an edge corresponds to the common part of two walls, one erected on $P_F \cap S$ (respectively, on $P_C \cap S$) and another erected on ∂T , for some $S, T \in \Gamma$. This part is composed of z -vertical segments that touch $P_F \cap S$ from above (respectively, $P_C \cap S$ from below), pass through ∂T , and touch P_C (respectively, P_F). We denote such edges by E_{21} to signify that they are formed by an intersection of two simplices, and by a boundary of a third simplex.

An interesting special case is the one in which $S = T$, i.e., one face corresponds to $P_F \cap S$ or $P_C \cap S$, and another to ∂S . Such edges correspond to a visibility wall erected on $P_F \cap \partial S$ or on $P_C \cap \partial S$, and are denoted by E_2 .

4. Edges incident to two blue faces. Such an edge corresponds to the common part of the visibility walls erected on ∂S and ∂T , for some $S, T \in \Gamma$. We denote such edges by E_{11} to signify that they are formed by two boundaries of simplices.

An interesting special case is the one in which $S = T$. In this case the edge is incident to two blue faces that correspond to walls erected on two incident two-dimensional features of ∂S , and it therefore corresponds to a visibility wall erected on a one-dimensional feature of ∂S . Such edges are denoted by E_1 .

To recap, the possible mnemonic representations of the edges of P_{3D} are E_3 , E_{22} , E_{21} , E_2 (a special case of E_{21}), E_{11} , and E_1 (a special case of E_{11}).

Each y -vertical visibility event inside P_{3D} corresponds to a y -vertical segment s that lies completely inside P_{3D} , such that s connects a point p on an edge of P_{3D} to a point p' on another edge of P_{3D} . Notice that p (as well as p') is a z -vertical projection of a specific z -vertical segment e (respectively, e'), whose bottom end-point lies on P_F , and whose top end-point lies on P_C . By construction, e and e' lie inside a common yz -parallel (two-dimensional) plane, which we denote by $\Pi_{e,e'}$. The two-dimensional feature of $\mathcal{V}_2(\Gamma)$ that corresponds to s (that is, the “wall” that is erected on s inside P in \mathbb{R}^4) is the trapezoid that has e and e' as its bases. This trapezoid is necessarily disjoint from Γ in its interior. We sometimes denote a y -vertical visibility event by (e, e') , where e and e' are as above, and where the points on e have a smaller y -coordinate than the points on e' .

Figure 1 provides an exhaustive visual catalogue of the possible types of y -vertical visibility events. There are 21 such types, determined by the type of the edge that contains the point p and the type of the edge that contains the point p' . For example, one such type of event is E_3E_{22} , in which the edge containing p is of type E_3 and the edge containing p' is of type E_{22} (or vice versa; this type can also be denoted by $E_{22}E_3$). For each type of event, the figure shows the (one or two) possible combinatorially distinct configurations of simplices inside the plane $\Pi_{e,e'}$ (up to the obvious symmetries of interchanging top and bottom and/or left and right); the cross section of the simplices inside $\Pi_{e,e'}$ is shown. The z -vertical segments e and e' are shown dashed. In each part of the figure, the trapezoid that is bounded by e , e' , P_F , and P_C , corresponding to a specific visibility event, can be clearly distinguished. By definition, the interior of this trapezoid has to be disjoint from all simplices of Γ . Notice that Fig. 1 treats E_2 and E_1 as distinct types of edges. They will also be treated as such in the analysis.

As shown in Fig. 1, some types of events have two possible combinatorially distinct configurations. To distinguish between them during the analysis, we sometimes describe an event not only by its type, but also by the configuration. For instance, we distinguish between $E_3E_3(\text{I})$ and $E_3E_3(\text{II})$ events, which are events of type E_3E_3 that have either the configuration shown in Fig. 1(b,I) or the one shown in Fig. 1(b,II).

3. Visibility Events in Arrangements of Hyperplanes

Analyzing vertical decompositions of arrangements of hyperplanes is simpler than in the more general case of 3-simplices, since hyperplanes do not have boundaries, and thus only three kinds of y -vertical visibility events can occur. Using the notation introduced in the previous section, these are E_3E_3 , E_3E_{22} , and $E_{22}E_{22}$ events. In order to prove the main result of this section, stated below, it is therefore sufficient to analyze only events of these three kinds.

Theorem 3.1. *The number of cells in the vertical decomposition of an arrangement of n hyperplanes in four dimensions is $O(n^4)$.*

Proof. In the following sequence of four lemmas, we analyze each kind of event in turn, and prove that there can only be $O(n^4)$ events of each kind. This is accomplished

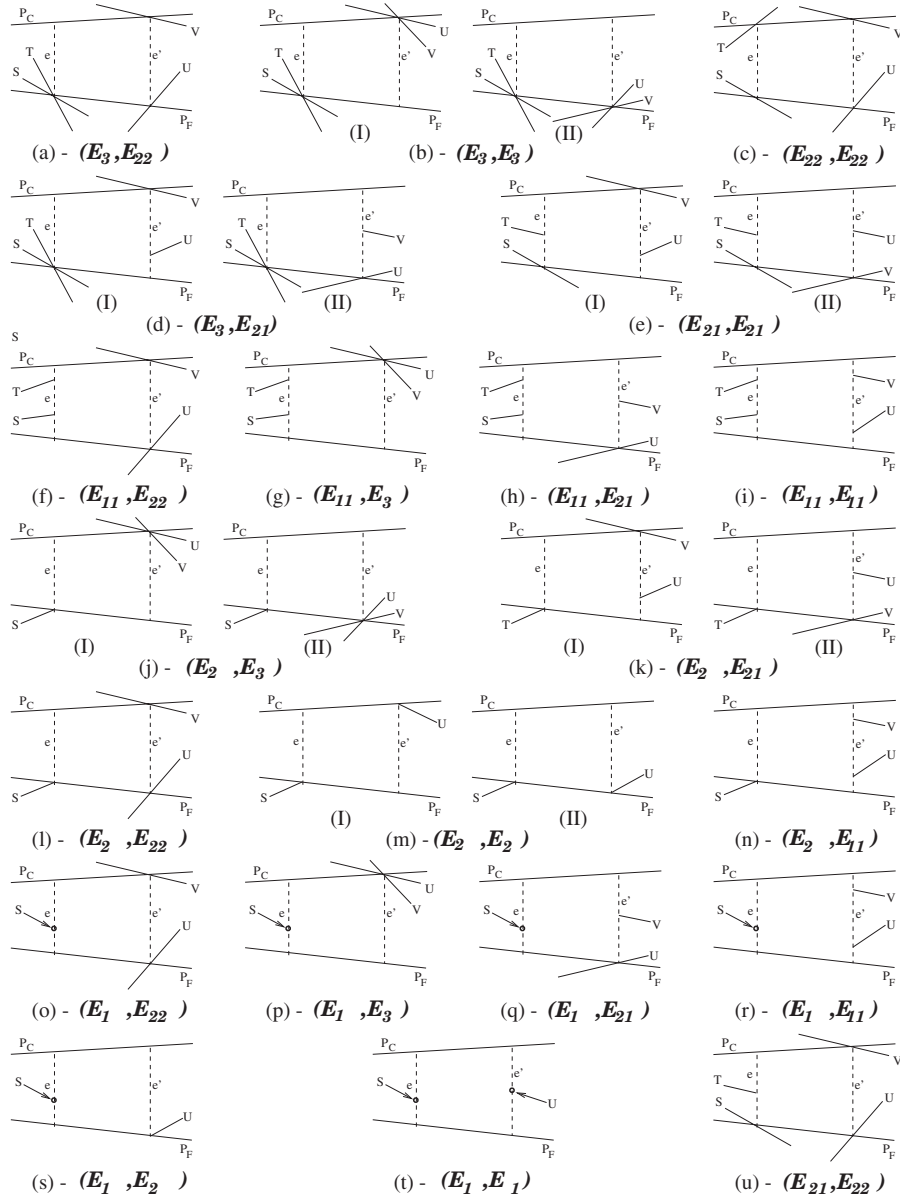


Fig. 1. All the possible types of y -vertical visibility events in arrangements of simplices (up to symmetry).

by charging each event to features of $\mathcal{A}(\Gamma)$ or of $\mathcal{V}_1(\Gamma)$, or to events that were analyzed previously. Special care is taken to ensure that each feature or event is only charged at most a constant number of times.

Lemma 3.2. *The number of $E_3E_3(\text{II})$ events (that is, events of type E_3E_3 that are in the configuration shown in Fig. 1(b,II)) in an arrangement of n hyperplanes in 4-space is $O(n^4)$.*

Proof. It is easy to see that each $E_3E_3(\text{II})$ event uniquely corresponds to a feature of the vertical decomposition of the three-dimensional arrangement $\mathcal{A}(\Gamma \cap P_F)$ inside the subspace P_F , where $P_F \in \Gamma$ is the hyperplane that contains the two edges that define the event. (The symmetric case, where the event occurs on the ceiling P_C of a cell, is treated in exactly the same manner.) Indeed, parametrize P_F by x, y, w and let y be the vertical direction. Consider the edges $\varphi \subseteq P_F \cap S \cap T$ and $\psi \subseteq P_F \cap U \cap V$ that define a particular $E_3E_3(\text{II})$ event and consider the planes $S^* = P_F \cap S, T^* = P_F \cap T, U^* = P_F \cap U, V^* = P_F \cap V$ of $\mathcal{A}(\Gamma \cap P_F)$. $\varphi \subseteq S^* \cap T^*$ and $\psi \subseteq U^* \cap V^*$ are edges of $\mathcal{A}(\Gamma \cap P_F)$ that are y -vertically visible in the subspace P_F . The considered $E_3E_3(\text{II})$ event uniquely corresponds to the y -vertical segment that connects φ and ψ in P_F and is a feature of the vertical decomposition of $\mathcal{A}(\Gamma \cap P_F)$. For any $P_F \in \Gamma$, the number of such features is $O(n^3)$ [2], [13]. Thus, the overall number of $E_3E_3(\text{II})$ events is $\sum_{P_F} O(n^3) = O(n^4)$. \square

From now on, throughout the remainder of this section, we only consider $E_3E_3(\text{I})$ events (as in Fig. 1(b,I)), and refer to them briefly as E_3E_3 events.

Lemma 3.3. *The number of E_3E_3 events in an arrangement of n hyperplanes in 4-space is $O(n^4)$.*

Proof. For any three hyperplanes $P_F, S, T \in \Gamma$, we slide a point a on the line $P_F \cap S \cap T$ at constant speed (in any of the two possible directions, say in the positive direction of the w -axis, from $-\infty$ to $+\infty$). Consider a two-dimensional yz -parallel plane Π_a that is attached to the point a , so that it contains a at all times during the sliding. The plane Π_a “sweeps” (a part of) $\mathcal{A}(\Gamma)$ in a fixed direction at a fixed speed. Thus, it contains a dynamic arrangement of lines, such that each line moves with a fixed speed, and the slope of each line is fixed. Each such moving line corresponds to a hyperplane; two moving lines intersect in a moving point, corresponding to an intersection of two hyperplanes, that also moves along a linear trajectory with a fixed speed. At discrete moments in time, during the sliding, three lines intersect in a point—such “critical events” correspond to intersections of three hyperplanes that are swept by Π_a . Three of the moving lines are always in degenerate configuration, as they intersect in a point at all times. These are the lines (that correspond to) P_F, S , and T . (Abusing the notation slightly, we denote $X \cap \Pi_a$, for $X \in \Gamma$, simply by X , and similarly for any other feature.)

As described in detail below, during this sliding we associate each E_3E_3 event that involves $P_F \cap S \cap T$ with a pair (U, Λ) , where $U \in \Gamma$ and Λ is an edge in the (four-dimensional) zone of U . Since we go over all triples (P_F, S, T) in this fashion, all E_3E_3 events in $\mathcal{V}_2(\Gamma)$ will be associated with such pairs at the end of the process. We will prove that each such pair can only be charged (i.e., associated with) a constant number of times, which will imply that the number of E_3E_3 events is asymptotically the same as the number of such pairs in the whole arrangement. The latter number is $O(n^4)$, since there are n hyperplanes U , and the zone of each contains $O(n^3)$ edges Λ [7]. We now describe the charging in detail.

Consider one E_3E_3 event (e, e') that involves $P_F \cap S \cap T$. Assume, without loss of generality, that the bottom end-point of the segment e lies on $P_F \cap S \cap T$, and the top end-point of e' lies on an intersection of three other hyperplanes of $\Gamma, P_C \cap U \cap V$, such that an upward z -vertical ray obtained by extending e upwards hits U before hitting V




Fig. 2. The types of E_3E_3 and E_3E_{22} events. (a), (b) E_3E_3 events of types A and B, respectively. (c), (d) E_3E_{22} events of types A and B. All illustrations show a representative configuration inside the plane Π_a at time $t_0 + \varepsilon$.

(see Fig. 1(b,I)). By definition, there is a specific moment in time (denote it by t_0) during the sliding, at which the plane Π_a coincides with the plane $\Pi_{e,e'}$. At time $t_0 + \varepsilon$, where $\varepsilon > 0$ is infinitesimally small, the y -coordinate of the point $P_C \cap V$ is either smaller or larger than the y -coordinate of the point $P_C \cap U$. These two possibilities distinguish between E_3E_3 events of type A and type B, respectively; see Fig. 2.

Suppose the event (e, e') is of type A. Notice that, immediately after time t_0 , U is “separated” (within Π_a) from $P_F \cap S \cap T$ by P_C and V , in the sense that U does not intersect the wedge bounded by P_C and V that contains $P_F \cap S \cap T$ at time t_0 . Due to linearity of the trajectories, this implies that no event of type E_3E_3 involving both $P_F \cap S \cap T$ and U can happen after time t_0 while $P_F \cap S \cap T$ lies inside this wedge. Notice that $P_F \cap S \cap T$ cannot cease to lie inside this wedge before intersecting either P_C or V . Such intersections correspond to vertices of $\mathcal{A}(\Gamma)$.

Let $\Lambda \subset P_F \cap S \cap T$ be the edge of $\mathcal{A}(\Gamma)$ that contains the bottom end-point of e . The preceding argument implies that no E_3E_3 event involving Λ and U can happen after time t_0 . We associate the event (e, e') with the pair (U, Λ) . Uniqueness of this association is ensured by the fact that, after the first such charge is made to some pair (U, Λ) (during sliding on the edge Λ), no other E_3E_3 event of type A involving U and Λ can occur.

Events (e, e') of type B are handled by a fully symmetric charging argument, in which the direction of the sliding of a is reversed.

We have thus shown that each E_3E_3 event can be associated with a pair (U, Λ) as described above, such that each such pair (U, Λ) is associated with at most one event of type A and at most one event of type B. This completes the proof of the lemma. \square

Lemma 3.4. *The number of E_3E_{22} events in an arrangement of n hyperplanes in 4-space is $O(n^4)$.*

Proof. For any $P_F, S, T \in \Gamma$, we slide a point a on the line $P_F \cap S \cap T$, with the yz -parallel plane Π_a attached to it, as in the proof of Lemma 3.3. During the sliding, each E_3E_{22} event (e, e') that involves $P_F \cap S \cap T$ is associated, as described below, with a pair (V, Λ) , where $V \in \Gamma$ and Λ is an edge in the zone of V . Since we go over all triples (P_F, S, T) in this fashion, all E_3E_{22} events in $\mathcal{V}_2(\Gamma)$ will be associated with such pairs at the end of the process. As shown in the proof of Lemma 3.3, the number of such pairs is $O(n^4)$. We will show that each such feature is only charged a constant number of times, which implies that the number of E_3E_{22} events is also $O(n^4)$.

Consider one E_3E_{22} event (e, e') , such that a coincides with the bottom end-point of e at time t_0 . (All E_3E_{22} events (e, e') , such that a coincides with the top end-point of e or

with an end-point of e' , can be handled analogously.) Suppose that the bottom end-point of e' lies on $P_F \cap U$ and the top one on $P_C \cap V$, for some $U, V \in \Gamma$. The y -coordinate of the point $P_F \cap U$ (inside the plane Π_a) at time $t_0 + \varepsilon$ is either smaller or larger than that of the point $P_C \cap V$. These two possibilities distinguish between E_3E_{22} events of type A and type B, respectively; see Fig. 2. Assume that the event (e, e') is of type A. The other case can be handled by a symmetric argument with the direction of the sliding reversed, as in the proof of Lemma 3.3.

Immediately after time t_0 , V is “separated” from $P_F \cap S \cap T$ by P_C and U , similarly to the proof of Lemma 3.3, such that no event of type E_3E_{22} involving both $P_F \cap S \cap T$ and V can happen after time t_0 before $P_F \cap S \cap T$ intersects either P_C or U . Such an intersection corresponds to a vertex of $\mathcal{A}(\Gamma)$.

Let $\Lambda \subset P_F \cap S \cap T$ be the edge of $\mathcal{A}(\Gamma)$ that contains the bottom end-point of e . We have just argued that no E_3E_{22} event involving Λ and V can happen after time t_0 . We can thus associate the event (e, e') with the pair (V, Λ) , and, as in the proof of Lemma 3.3, this association is unique. This completes the proof of the lemma. \square

Remark 1. The proof of Lemma 3.4 can be modified in a straightforward fashion to bound by $O(n^4)$ the number of events of a more general type, which we call $G_{3,22}$. A $G_{3,22}$ event is said to happen between three points, $p_1 \in P_F \cap S \cap T$, $p_2 \in P_F \cap U$, and $p_3 \in P_C \cap V$, for some $P_F, P_C, S, T, U, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, the open segment (p_1, p_2) is disjoint from all hyperplanes of Γ other than P_F , the open segment (p_2, p_3) is z -vertical and is disjoint from all hyperplanes of Γ , and, furthermore, S, U , and V lie outside the trapezoid that has as its four vertices p_1, p_2, p_3 , and the point on P_C that is covertical with p_1 ; see Fig. 3.

Lemma 3.5. *The number of $E_{22}E_{22}$ events in an arrangement of n hyperplanes in 4-space is $O(n^4)$.*

Proof. For any $P_F, P_C, S, T \in \Gamma$, we slide a z -vertical segment a , such that its bottom end-point always lies on $P_F \cap S$ and its top end-point always lies on $P_C \cap T$, with the yz -parallel plane Π_a attached to it, similarly to the proof of Lemma 3.3. Assuming general position, the locus of points on $P_C \cap T$ that are intersected by the z -vertical hyperplane spanned by $P_F \cap S$ is a line. This implies that the trajectory of a is linear and one-dimensional, and therefore Π_a contains a dynamic arrangement of lines, as in the proof of Lemma 3.3.




Fig. 3. A $G_{3,22}$ event. Notice that, as opposed to E_3E_{22} events, in the case of $G_{3,22}$ events the trapezoid that has e and e' as its bases is not necessarily empty.




Fig. 4. The analysis of $E_{22}E_{22}$ events. (a) The configuration inside Π_a at time $t_0 + \varepsilon$; (b)–(d) the three types of events that are associated with $E_{22}E_{22}$ events and occur at a later time.

During the sliding, each $E_{22}E_{22}$ event (e, e') that involves a is associated, as described below, with either a feature of $\mathcal{V}_1(\Gamma)$ or a $G_{3,22}$ event. Since we go over all quadruples (P_F, P_C, S, T) in this fashion, all $E_{22}E_{22}$ events in $\mathcal{V}_2(\Gamma)$ will be associated with such features at the end of the process. The complexity of $\mathcal{V}_1(\Gamma)$ is $O(n^4)$, and Remark 1 states that the number of $G_{3,22}$ events is also $O(n^4)$. As we will show, each such feature is only charged a constant number of times, which implies that the number of $E_{22}E_{22}$ events is also $O(n^4)$.

Consider one E_3E_{22} event (e, e') , such that a coincides with e at time t_0 . (All $E_{22}E_{22}$ events (e, e') , such that a coincides with e' , can be handled analogously.) Suppose that the bottom end-point of e' lies on $P_F \cap U$ and the top one on $P_C \cap V$, for some $U, V \in \Gamma$. Assume, without loss of generality, that the y -coordinate of the point $P_F \cap U$ (inside the plane Π_a) at time $t_0 + \varepsilon$ is smaller than the y -coordinate of the point $P_C \cap V$; see Fig. 4(a).

Claim. No $E_{22}E_{22}$ event (g, g') can occur, such that the segment a coincides with g , after time t_0 and before some (moving) line X (corresponding to a hyperplane $X \in \Gamma$, which may be U) intersects either (1) the top end-point of a on $P_C \cap T$ (Fig. 4(b)), or (2) the bottom end-point of a on $P_F \cap S$ (Fig. 4(c)), or (3) $P_C \cap V$ (Fig. 4(d)).

Proof. Suppose to the contrary that such an event (g, g') does occur before any of these intersections. After time t_0 and before any event of the three types listed in the claim happens, the relative interior of the segment connecting $P_C \cap T$ and $P_C \cap V$ is disjoint from all hyperplanes of Γ other than P_C . This implies that the top end-point of g' has to lie on $P_C \cap V$ (while the bottom end-point has to lie on P_F). However, in this case g' is necessarily intersected by U . This prevents (g, g') from being an $E_{22}E_{22}$ event, since the segment g' has to be disjoint in its interior from all the hyperplanes of Γ if (g, g') is a y -vertical visibility event. \square

The claim implies that we can associate the event (e, e') with the first time (after time t_0) some line X causes one of the three events listed in the claim to happen, if at least one of these events does occur during the sliding. Notice that an event of type (3) is a $G_{3,22}$ event, while the first (in time) event of type (1) or (2) corresponds to a feature of $\mathcal{V}_1(\Gamma)$. As in the proof of Lemma 3.3, each of these events can be charged at most twice, once for each direction of the sliding of a . This implies that the number of $E_{22}E_{22}$ events that can be associated in this fashion is $O(n^4)$. The only case in which such association cannot be made is when the segment a continues sliding ad infinitum without any event of type (1), (2), or (3) occurring. In this case the claim implies that no $E_{22}E_{22}$ events occur during the sliding as well. We can thus simply charge the event (e, e') to the quadruple

(P_F, P_C, S, T) . The number of such quadruples is $O(n^4)$, and each quadruple can be charged at most twice. This completes the proof of the lemma. \square

Lemmas 3.2–3.5 bound the number of all y -vertical visibility events by $O(n^4)$, thereby implying Theorem 3.1. \square

4. Visibility Events in Arrangements of Simplices

The analysis in the case of simplices is substantially more involved than in the case of hyperplanes, since all of the 21 types of events illustrated in Fig. 1 can occur. In this case we have:

Theorem 4.1. *The number of cells in the vertical decomposition of an arrangement of n 3-simplices in four dimensions is $O(n^4\alpha(n)\log^2 n)$.*

Overview of the Proof. The theorem is proved in two major stages. In the first stage, given in Sections 4.1 and 4.2, we analyze all types of events, with the exception of E_3E_3 events, and prove that the number of events of each type is $O(n^4\alpha(n)\log n)$. This stage often uses ideas similar to those introduced in the previous section. That is, events are charged to features of $\mathcal{A}(\Gamma)$ and $\mathcal{V}_1(\Gamma)$, and to events of types that were analyzed previously. Special care is taken to ensure that such features or events are only charged at most a constant number of times.

In the second stage of the proof, described in Section 4.3, we bound the number of E_3E_3 events, using the probabilistic technique introduced by Tagansky for analyzing substructures in arrangements of linear surfaces [13]. A charging scheme is presented, such that each E_3E_3 event is charged either to a feature of $\mathcal{A}(\Gamma)$, or to a feature of $\mathcal{V}_1(\Gamma)$, or to an event of one of the types that were analyzed in the first stage, or to a 1-level E_3E_3 event (see Section 4.3 for definitions). To utilize the Tagansky technique, we carefully bound the number of times each 1-level E_3E_3 event is charged in this fashion. In particular, we prove that although each E_3E_3 event makes 4 “units” of charge, each 1-level E_3E_3 event is charged by at most 2 such units. As a result, we obtain a recurrence for the number of E_3E_3 events, which does solve to the asserted bound of $O(n^4\alpha(n)\log^2 n)$, thus completing the proof.

Proof of Theorem 4.1. The proof is laid out in the following three subsections.

4.1. Reductions to Lower Envelopes

Lemma 4.2. *The number of $E_1E_1, E_1E_2, E_2E_2, E_1E_{11}, E_1E_3, E_1E_{21}, E_{11}E_{11}, E_{11}E_2,$ and $E_2E_3(\text{II})$ events in an arrangement of n 3-simplices in 4-space is $O(n^4)$.*

Proof. Each event of these types is defined by an interaction of at most four simplices. Any quadruple of simplices of Γ can define $O(1)$ such events, and there are $O(n^4)$ such quadruples. Therefore, the number of such events is $O(n^4)$. For example, the $E_{11}E_{11}$

event depicted in Fig. 1(i) is defined by the simplices S, T, U , and V . That is, a fixed quadruple $S, T, U, V \in \Gamma$ defines $O(1)$ configurations in \mathbb{R}^4 , each of which can be an $E_{11}E_{11}$ event. In each such configuration, the simplices P_F and P_C are fixed automatically. A similar property holds for all the types of events that are listed in the lemma. \square

In what follows we introduce a number of *generalized* types of events, starting with $G_{11,1}$, $G_{11,2}$, and $G_{21,1}$ in the following definition. These types will be helpful in subsequent sections, and this is the main motivation behind their introduction.

Definition.

- A $G_{11,1}$ event is said to happen between three points, $p_1 \in \partial S$, $p_2 \in \partial T$, and $p_3 \in \partial U$, for some $S, T, U \in \Gamma$, if all three points lie inside a common yz -parallel plane, the segment (p_1, p_2) is z -vertical, and the point p_3 is contained in a one-dimensional edge of U .
- A $G_{11,2}$ event is said to happen between three points, $p_1 \in \partial S$, $p_2 \in \partial T$, and $p_3 \in C \cap \partial U$, for some $C, S, T, U \in \Gamma$, if all three points lie inside a common yz -parallel plane and the segment (p_1, p_2) is z -vertical.
- A $G_{21,1}$ event is said to happen between three points, $p_1 \in F \cap S$, $p_2 \in \partial T$, and $p_3 \in \partial U$, for some $F, S, T, U \in \Gamma$, if all three points lie inside a common yz -parallel plane, the segment (p_1, p_2) is z -vertical, and the point p_3 is contained in a one-dimensional edge of U .

It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4)$ as well.

Lemma 4.3. *The number of $E_2E_3(I)$ events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n))$.*

Proof. For every three simplices $S, P_F, P_C \in \Gamma$, consider all the $E_2E_3(I)$ events (e, e') , in which the bottom end-point of e lies on $\partial S \cap P_F$, and the top end-point of e lies on P_C (see Fig. 1(j,I)). (Other events involving S, P_F , and P_C , such as events (e, e') , in which the top end-point of e lies on $\partial S \cap P_F$, and the bottom end-point of e lies on P_C can be treated in a completely analogous fashion.) We show below that the number of such events is $O(n\alpha(n))$. Since there are $O(n^3)$ such triples (S, P_F, P_C) , this will imply that the overall number of $E_2E_3(I)$ events is $O(n^4\alpha(n))$.

Parametrize the hyperplane P_C in terms of the coordinates x, y , and w , such that a point (x_0, y_0, w_0) in this parametrization corresponds to the unique point on P_C that is intersected by the z -vertical line spanned by the point $(x_0, y_0, 0, w_0) \in \mathbb{R}^4$. Consider the collection of polygons $\Gamma' = \{\gamma \cap P_C \mid \gamma \in \Gamma\}$, parametrized as above. In this parametrization, $\mathcal{A}(\Gamma')$ is an arrangement of n convex polygons in 3-space, where each nonempty polygon is of constant complexity.

The locus of points that lie on P_C , and are intersected by the infinite z -vertical wall spanned by $\partial S \cap P_F$ is a collection of $O(1)$ line segments. Consider one of these segments, denoted by f , in the above parametrization of P_C . Consider the y -vertical lower envelope of the part of $\mathcal{A}(\Gamma')$ that lies y -vertically above this segment f (in this parametrization).

By definition, each event (e, e') as above uniquely corresponds to a vertex of such a lower envelope, for one of the segments f . Since this is a lower envelope of $O(n)$ line segments, it has complexity $O(n\alpha(n))$ [1], and the lemma follows. \square

Definition. A $G_{2,3}$ event is said to happen between three points, $p_1 \in \partial S \cap F$, $p_2 \in C$, and $p_3 \in C \cap U \cap V$, for some $F, C, S, U, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, the segment (p_1, p_2) is z -vertical, and the open segment (p_2, p_3) is disjoint from all simplices of Γ other than C . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n))$ as well.

Lemma 4.4. *The number of $E_{11}E_3$ events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n))$.*

Proof. This proof is very similar to the proof of Lemma 4.3. For every three simplices $S, T, P_C \in \Gamma$, consider all the $E_{11}E_3$ events (e, e') , in which e touches ∂S and ∂T , and the top end-point of e lies on P_C (see Fig. 1(g); as above, symmetric cases are handled analogously). We show below that the number of such events is $O(n\alpha(n))$, which implies that the overall number of $E_{11}E_3$ events is $O(n^4\alpha(n))$.

Parametrize P_C as in the proof of Lemma 4.3. The locus of points that lie on P_C and are intersected by the infinite z -vertical wall spanned by ∂S and ∂T , is a collection of $O(1)$ line segments. As in the proof of Lemma 4.3, each event (e, e') corresponds to a vertex of the y -vertical lower envelope of the part of $\mathcal{A}(\Gamma')$ that lies y -vertically above one of these segments (in this parametrization of P_C), which implies, as above, the stated bound. \square

Definition. A $G_{11,3}$ event is said to happen between three points, $p_1 \in \partial S$, $p_2 \in \partial T$, and $p_3 \in C \cap U \cap V$, for some $C, S, T, U, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, the segment (p_1, p_2) is z -vertical, and the segment that connects p_3 to the point on C that lies z -vertically above/below the point p_1 is disjoint from all simplices of Γ other than C . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n))$ as well.

Lemma 4.5. *The number of E_1E_{22} events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n))$.*

Proof. This proof is similar to the proof of Lemma 4.3. For every three simplices $S, P_F, U \in \Gamma$, consider all the E_1E_{22} events (e, e') , in which e lies on a one-dimensional edge φ of S and the bottom end-point of e' lies on $P_F \cap U$ (see Fig. 1(o); symmetric cases are handled analogously). We show below that the number of such events is $O(n\alpha(n))$, which implies that the overall number of E_1E_{22} events is $O(n^4\alpha(n))$.

The locus of points on $P_F \cap U$ that lie in a common yz -parallel plane with a point on φ is a line segment. Consider the z -vertical lower envelope, denoted by L , of the part of $\mathcal{A}(\Gamma)$ that lies above this line segment (in the z -direction). It is easy to see that all events (e, e') as above uniquely correspond to vertices of L , which is a lower envelope of an arrangement of $O(n)$ line segments in a plane. The number of these vertices is $O(n\alpha(n))$, and the lemma follows. \square

Definition. A $G_{1,22}$ event is said to happen between three points, $p_1 \in F \cap S$, $p_2 \in C \cap T$, and $p_3 \in \partial U$, for some $F, C, S, T, U \in \Gamma$, if all three points lie inside a common yz -parallel plane, the open segment (p_1, p_2) is z -vertical and is disjoint in its interior from all simplices of Γ , and the point p_3 is contained in a one-dimensional edge of U . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n))$ as well.

Lemma 4.6. *The number of events of types E_3E_{22} , E_2E_{22} , and $E_3E_{21}(\text{II})$ in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n) \log n)$.*

Proof. For every two simplices $P_F, U \in \Gamma$, consider all events (e, e') of types E_3E_{22} , E_2E_{22} , and $E_3E_{21}(\text{II})$, such that the bottom end-point of e' lies on $P_F \cap U$ (as in Figs. 1(a), 1(l), and 1(d,II); symmetric cases are handled analogously). We show below that the number of such events is $O(n^2\alpha(n))$. Since there are $O(n^2)$ such pairs (P_F, U) , this will imply that the overall number of E_3E_{22} , E_2E_{22} , and $E_3E_{21}(\text{II})$ events is $O(n^4\alpha(n))$.

Consider the z -vertical lower envelope, denoted by L_1 , of the part of $\mathcal{A}(\Gamma)$ that lies above $P_F \cap U$ (in the z -direction). This is a lower envelope of at most n planar convex polygons, each of constant complexity, within the three-dimensional z -vertical wall spanned by $P_F \cap U$.

Parametrize P_F as in the proof of Lemma 4.3, and consider the collection of polygons $\Gamma' = \{\gamma \cap P_F \mid \gamma \in \Gamma\}$, in this parametrization. $\mathcal{A}(\Gamma')$ is an arrangement of n convex polygons in 3-space, each polygon being of constant complexity. Consider the polygon of Γ' that corresponds to $U \cap P_F$, and consider the y -vertical upper envelope, denoted by L_2 , of the part of $\mathcal{A}(\Gamma')$ that lies below this polygon (in the y -direction), with respect to the above parametrization.

Consider an overlay $O(L_1, L_2)$ of L_1 and L_2 , such that a point $q \in L_1$ is overlaid with a point $q^* \in L_2$ if the z -vertical projection of q onto $P_F \cap U$ coincides with the y -vertical projection of q^* onto $P_F \cap U$. Every point in $O(L_1, L_2)$ thus corresponds to a z -vertical ray (in 4-space) and a y -vertical ray (along P_F) emanating from a common point on $P_F \cap U$. By definition, each event (e, e') as above corresponds to a vertex of $O(L_1, L_2)$, at which a z -projected edge of L_1 crosses a y -projected edge of L_2 . As shown in [11], the complexity of each such overlay is $O(n^2\alpha(n) \log n)$, and this implies the lemma. \square

Definition. A $G_{3,22}$ (resp., $G_{2,22}$, $G_{3,21}(\text{II})$) event is said to happen between three points, $p_1 \in F \cap S \cap T$ (resp., $p_1 \in F \cap \partial S$ or $p_1 \in \partial F \cap S$, $p_1 \in F \cap S \cap T$), $p_2 \in F \cap U$ or $p_2 \in \partial F$, and $p_3 \in C \cap V$ (resp., $p_3 \in C \cap V$, $p_3 \in \partial V$), for some $F, C, S, T, U, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, the open segment (p_1, p_2) is disjoint from all simplices of Γ other than F , and the open segment (p_2, p_3) is z -vertical and is also disjoint from all simplices of Γ . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n) \log n)$ as well.

Lemma 4.7. *The number of $E_{21}E_{21}$, $E_{21}E_2$, and $E_{21}E_{11}$ events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n) \log n)$.*

Proof. For every two simplices $T, U \in \Gamma$, consider all the $E_{21}E_{21}$, $E_{21}E_2$, and $E_{21}E_{11}$

events (e, e') , such that e touches ∂T and e' touches ∂U . We will show that the number of such events is $O(n^2\alpha(n) \log n)$, and, as in the proof of Lemma 4.6, this will imply the lemma.

Consider the z -vertical lower envelope, denoted by L_1 , of the part of $\mathcal{A}(\Gamma)$ that lies above ∂T in the z -direction. Similarly, consider the z -vertical upper envelope, denoted by U_1 , of the part of $\mathcal{A}(\Gamma)$ that lies below ∂T . Consider also L_2 and U_2 , which are defined in a similar fashion with respect to ∂U (instead of ∂T). Consider an overlay $O(L_1, L_2)$, such that a point $q \in L_1$ is overlaid with a point $q^* \in L_2$ if they lie in the same yz -parallel plane. The overlays $O(L_1, U_2)$, $O(U_1, L_2)$, and $O(U_1, U_2)$ are defined in a similar fashion to $O(L_1, L_2)$.

By definition, each event (e, e') as above corresponds to a vertex of $O(L_1, L_2)$, $O(L_1, U_2)$, $O(U_1, L_2)$, or $O(U_1, U_2)$. At such a vertex, a z -projected edge of L_1 or U_1 crosses a z -projected edge of L_2 or U_2 . (Concerning $E_{21}E_2$ events, shown in Fig. 1(k), notice that the intersection $P_F \cap \partial T$, for any $P_F \in \Gamma$, is by definition an edge of L_1 and of U_1 .) The complexity of these two-dimensional overlays is $O(n^2\alpha(n) \log n)$ [11], and the lemma follows. \square

Definition. A $G_{21,2}$ event is said to happen between three points, $p_1 \in F \cap \partial T$, $p_2 \in \partial U$, and $p_3 \in C \cap V$ or $p_3 \in F \cap V$, for some $F, C, T, U, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, and the open segment (p_2, p_3) is z -vertical and is disjoint from all simplices of Γ . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n) \log n)$ as well.

4.2. Charging Schemes

Lemma 4.8. *The number of E_3E_{21} events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n) \log n)$.*

Proof. The number of $E_3E_{21}(\text{II})$ events was bounded in Lemma 4.6. We thus only need to consider $E_3E_{21}(\text{I})$ events.

For any $P_F, P_C, S, T \in \Gamma$, we slide a z -vertical segment a , such that its bottom end-point lies on $P_F \cap S \cap T$ and its top end-point lies on P_C , at constant speed, in any of the two possible directions.

Similarly to the proof of Lemma 3.5, consider a two-dimensional yz -parallel plane Π_a that is attached to the segment a , so that Π_a contains a at all times during the sliding. The plane Π_a “sweeps” (a part of) $\mathcal{A}(\Gamma)$ in a fixed direction at a fixed speed. Thus it contains a dynamic arrangement of segments, such that the slope of each segment is fixed. Each such moving segment corresponds to a simplex of Γ .

During the sliding, each $E_3E_{21}(\text{I})$ event (e, e') that involves a will be associated with either a vertex of $\mathcal{A}(\Gamma)$, or an event of type $G_{11,3}$, $G_{2,3}$, or $G_{3,21}(\text{II})$. Since we go over all quadruples (P_F, P_C, S, T) in this fashion, all $E_3E_{21}(\text{I})$ events in $\mathcal{V}_2(\Gamma)$ will be associated with such features at the end of the process. As we will show, each such feature can only be charged a constant number of times, which implies that the number of $E_3E_{21}(\text{I})$ events is $O(n^4\alpha(n))$.

Consider one $E_3E_{21}(\text{I})$ event (e, e') , such that a coincides with e at time t_0 . (All $E_3E_{21}(\text{I})$ events (e, e') , such that a coincides with e' , can be handled analogously.)

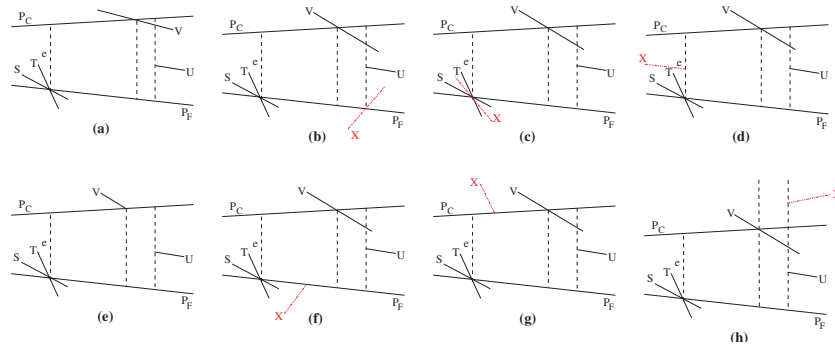


Fig. 5. The analysis of $E_3E_{21}(I)$ events. (a) Configuration inside Π_a at time $t_0 + \varepsilon$; (b)–(h) types of events that are associated with $E_3E_{21}(I)$ events and occur during the sliding.

Suppose that the top end-point of e' lies on $P_C \cap V$, the bottom end-point of e' lies on P_F , and e' touches ∂U in its interior, for some $U, V \in \Gamma$.

The y -coordinate of the point $P_C \cap V$ (inside the plane Π_a) at time $t_0 + \varepsilon$ is either smaller or larger than the y -coordinate of the point ∂U . These two possibilities distinguish between $E_3E_{21}(I)$ events of type A and type B, respectively. Assume that the event (e, e') is of type A (see Fig. 5(a)). The other case can be handled by a symmetric argument with the direction of the sliding reversed.

Claim. *No $E_3E_{21}(I)$ event (g, g') can occur, such that the segment a coincides with g , after time t_0 and before one of the following events happens within Π_a :*

- (1) ∂P_F or $X \cap P_F$, for some $X \in \Gamma$, passes below ∂U (see Fig. 5(b) for an illustration of the latter case).
- (2) A segment X intersects the bottom end-point of a (Fig. 5(c)).
- (3) An end-point of a segment X lies on the interior of a (Fig. 5(d)).
- (4) An end-point of V lies on P_C (Fig. 5(e)).
- (5) An end-point of a segment X lies on P_F (Fig. 5(f)).
- (6) An end-point of a segment X lies on the segment that connects the top end-point of a to $P_C \cap V$ (Fig. 5(g)).
- (7) An end-point of a segment X (which may be P_C) lies on the z -vertical line spanned by ∂U (Fig. 5(h)).
- (8) A one-dimensional feature (edge) of a simplex $X \in \Gamma$ intersects Π_a . (This happens, for example, when X first appears on the plane Π_a . At the moment such an event occurs, $X \cap \Pi_a$ is a single point. Another special case is when $X = U$, and $U \cap \Pi_a$ turns into a point and disappears.)

Proof. We assume that at some specific point of time after time t_0 none of the events specified in the claim has happened yet, and prove that this necessarily means that, at this point of time, no event (g, g') as above has occurred either. Note that after time t_0 and before one of the specified events occurs, all the trajectories inside Π_a are linear.

Consider the (dynamic) infinite z -vertical slab within Π_a , bounded from the right by the z -vertical line spanned by ∂U , and bounded from the left by the z -vertical line spanned by $P_C \cap V$. At time $t_0 + \varepsilon$, the width of this slab is infinitesimal and it contains

no ∂X in its interior, for any $X \in \Gamma$. As long as an event of type (4), (7), or (8) does not happen, the slab is well defined (that is, nonempty $\partial U \cap \Pi_a$ and $P_C \cap V \cap \Pi_a$ exist). Moreover, due to linearity of the trajectories, its width grows continuously with time. As long as an event of type (7) does not happen, no end-point of any segment intersects the right boundary of the slab, and as long as an event of type (8) does not happen, no new feature suddenly appears in the interior of the slab.

These observations together with the above assumption that none of the specified events have occurred imply that no end-point of any segment can intersect the left boundary of the slab from inside the slab (that is, intersect this left boundary at some time $t_1 > t_0$, after lying in the interior of the slab at time $t_1 - \varepsilon$). This is because we have essentially showed that any end-point of a segment that lies inside the slab could not have entered the slab by intersecting its right boundary, and could not have appeared in the interior of the slab without intersecting its boundary at all. Any end-point that lies inside the slab must have thus entered the slab by intersecting its left boundary, which means, due to linearity of all the trajectories and the fact that the slab continuously expands, that no end-point that lies inside the slab can exit the slab by intersecting the same left boundary.

Consider the trapezoid Δ bounded by a , P_F , P_C , and the left boundary of the slab. At time $t_0 + \varepsilon$, no end-point of any segment lies in its relative interior. Before an event of type (8) happens, no new feature suddenly appears inside it. Before an event of type (3) happens, no end-point penetrates Δ through its left edge. Before an event of type (5) or (6) happens, its top and bottom edges are also not penetrated by any segment end-point. We have showed above that its right edge (that lies on the left boundary of the slab discussed above) cannot be penetrated by an end-point from outside the trapezoid (and inside the slab). So no segment end-point can penetrate the boundaries of this trapezoid.

It is clear that in an $E_3E_{21}(I)$ event (g, g') as specified in the claim, g' either lies inside Δ or coincides with its right z -vertical edge. By definition, when the event (g, g') materializes, an end-point of some segment has to intersect g' . Due to the linearity of the trajectory of this end-point, immediately before the event it has to lie either to the right of g' (and thus lie inside the slab or in the interior of Δ) or to the left (and thus lie in the interior of Δ). Both options are impossible, since we have shown above that no end-point can intersect the left boundary of the slab from the right, and that no end-point can lie inside the trapezoid. Thus, our initial assumption indeed implies that no $E_3E_{21}(I)$ event (g, g') as specified in the claim can occur. \square

The claim implies that we can associate the event (e, e') with the first time (after time t_0) one of the events that are specified in the claim happens, if any of them do happen during the sliding. After time t_0 and before any of the above events happens, the relative interior of the segment connecting $P_F \cap S \cap T$ and the point on P_F that lies below ∂U is disjoint from all simplices of Γ , other than P_F .

This ensures the following. An event of type (2) corresponds to a vertex of $\mathcal{A}(\Gamma)$. An event of type (4) or (6) is a $G_{3,2}$ event. Notice also that an event of type (3), (5), or (8) is defined by an interaction of four simplices, as in Lemma 4.2, and thus there are $O(n^4)$ such events. An event of type (1) is a $G_{3,21}(II)$ event, and an event of type (7) is a $G_{11,3}$ event. (The latter is ensured, in part because events of type (1) do not occur before events of type (7); more accurately, whenever an event of type (1) occurs, it is immediately charged and the sliding is stopped.)

Since each of these events can only be charged a constant number of times, the number of $E_3E_{21}(\mathbb{I})$ events that can be associated with them as above is $O(n^4\alpha(n)\log n)$. The only case in which such an association cannot be made is when none of the above events occurs throughout the rest of the sliding, after time t_0 . In this case the claim implies that no $E_3E_{21}(\mathbb{I})$ event occurs during the sliding as well. We can thus simply charge the event (e, e') to the quadruple (P_F, P_C, S, T) . The number of such quadruples is $O(n^4)$, and the charging is unique up to the two possible directions of the sliding. This completes the proof of the lemma. \square

Definition. A $G_{3,21}(\mathbb{I})$ event is said to happen between three points, $p_1 \in C \cap V$, $p_2 \in F$, and $p_3 \in F \cap S \cap T$, for some $F, C, S, T, V \in \Gamma$, if all three points lie inside a common yz -parallel plane, the segment (p_1, p_2) is z -vertical and intersects ∂U , for some $U \in \Gamma$, and the interior of the trapezoid bounded by F, C , the z -vertical line spanned by p_1 , and the z -vertical line spanned by p_2 , does not meet any simplex of Γ other than U . It is easy to see that the above proof can be applied to bounding the number of such events by $O(n^4\alpha(n)\log n)$ as well.

Lemma 4.9. *The number of $E_{11}E_{22}$ events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n))$.*

Proof. We use a charging scheme similar to the one used to bound the number of $E_{22}E_{22}$ events in the case of hyperplanes (see Lemma 3.5).

For any $P_F, P_C, S, T \in \Gamma$, the locus of points on P_C that are intersected by the $O(1)$ z -vertical walls spanned by z -vertical lines that pass simultaneously through ∂S and ∂T , is a collection of $O(1)$ line segments. We treat one of them, denoted by f ; the others can be treated analogously. We slide a z -vertical segment a , such that its bottom end-point always lies on P_F and its top end-point slides on f , at constant speed (in any of the two possible directions), with the yz -parallel plane Π_a attached to it.

During the sliding, each (e, e') event that involves a is associated with either a vertex of $\mathcal{A}(\Gamma)$ or a feature of $\mathcal{V}_1(\Gamma)$ or an event of type $G_{11,1}$, $G_{11,2}$, or $G_{11,3}$. Since we go over all quadruples (P_F, P_C, S, T) in this fashion, all $E_{11}E_{22}$ events in $\mathcal{V}_2(\Gamma)$ will be associated with such features at the end of the process. As we will show, each such feature can only be charged (i.e., associated with) a constant number of times, which implies that the number of $E_{11}E_{22}$ events is $O(n^4\alpha(n))$.

Consider one $E_{11}E_{22}$ event (e, e') that involves a . By definition, there is a specific moment in time (denote it by t_0) during the sliding, at which the plane Π_a coincides with the plane $\Pi_{e,e'}$. Assume, without loss of generality, that the y -coordinate of the point $P_F \cap U$ at time $t_0 + \varepsilon$ is smaller than the y -coordinate of the point $P_C \cap V$ (as in Fig. 6).

Claim. *No $E_{11}E_{22}$ event (g, g') can occur such that the segment a coincides with g after time t_0 and before one of the following events happens within Π_a :*

1. A (moving) segment X (corresponding to some simplex $X \in \Gamma$) intersects the top end-point of a (see Fig. 6(a)).
2. A segment X (which may be U) intersects the bottom end-point of a (Fig. 6(b)).
3. An end-point of a segment X meets the interior of a (Fig. 6(c)).




Fig. 6. The various types of events that are associated with $E_{11}E_{22}$ events and occur during the sliding.

4. An end-point of U intersects P_F or an end-point of P_F intersects U (see Fig. 6(d) for an illustration of the former case).
5. An end-point of V intersects P_C or an end-point of P_C intersects V (see Fig. 6(e) for an illustration of the former case).
6. An end-point of a segment X intersects the segment that connects $P_F \cap U$ with the bottom end-point of a (Fig. 6(f)).
7. An end-point of a segment X intersects the segment that connects $P_C \cap V$ with the top end-point of a (Fig. 6(g)).
8. A segment X (which may be U) intersects $P_C \cap V$ (Fig. 6(h)).
9. A one-dimensional feature (edge) of a simplex $X \in \Gamma$ intersects Π_a .

Proof. After time t_0 and before one of the above events happens, all the trajectories are linear, and the segment that connects $P_C \cap V$ with the top end-point of a is disjoint in its interior from all simplices of Γ other than P_C . Therefore, in any $E_{11}E_{22}$ event (g, g') as specified in the claim, the top end-point of g' has to lie on $P_C \cap V$ (while the bottom end-point has to lie on P_F). However, in this case, the trapezoid (g, g') is necessarily intersected by U , which prevents (g, g') from being an $E_{11}E_{22}$ event. \square

The claim implies that we can associate the event (e, e') with the first time (after time t_0) one of the events specified in the claim happens, if any of them occurs at all. Any event of types (4)–(7) is a $G_{11,2}$ event, while an event of type (8) is a $G_{11,3}$ event, and an event of type (9) is a $G_{11,1}$ event. Notice that after time t_0 and prior to any of these events, the segment a is disjoint in its interior from all simplices of Γ (aside from the end-point of S and T). This ensures that any event of types (1)–(3) corresponds to a feature of $\mathcal{V}_1(\Gamma)$. If none of the events specified in the claim happens, we can charge the event (e, e') to the quadruple (P_F, P_C, S, T) , as in the proof of Lemma 4.8. This completes the proof of the lemma. \square

Lemma 4.10. *The number of $E_{21}E_{22}$ events in an arrangement of n 3-simplices in 4-space is $O(n^4 \alpha(n) \log n)$.*

Proof. We use a charging scheme similar to the one used in Lemma 4.9. For any $P_F, P_C, S, T \in \Gamma$, the locus of points on P_C that are intersected by the $O(1)$ z -vertical

walls spanned by z -vertical lines that pass simultaneously through ∂T and $P_F \cap S$ is a collection of $O(1)$ line segments. We slide a z -vertical segment a , such that its bottom end-point always lies on $P_F \cap S$, it touches ∂T in its interior, and its top end-point slides on one of these segments on P_C , at constant speed.

Similarly to the proof of Lemma 4.9, consider a two-dimensional yz -parallel plane Π_a that is attached to segment a , so that Π_a contains a at all times during the sliding. In exactly the same manner as in Lemma 4.9, we can associate each event (e, e') that involves a with either a feature of $\mathcal{V}_1(\Gamma)$ or an event of type $G_{21,1}$, $G_{21,2}$, or $G_{3,21}(\text{II})$. Since we go over all quadruples (P_F, P_C, S, T) in this fashion, all $E_{21}E_{22}$ events in $\mathcal{V}_2(\Gamma)$ will be associated with such features at the end of the process. Since each such feature can only be charged a constant number of times, this implies that the number of $E_{21}E_{22}$ events is $O(n^4\alpha(n) \log n)$. \square

Lemma 4.11. *The number of $E_{22}E_{22}$ events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n) \log n)$.*

Proof. We use a charging scheme similar to the one used in Lemma 4.9. For any $P_F, P_C, S, T \in \Gamma$, the locus of points on $P_C \cap T$ that are intersected by the $O(1)$ z -vertical walls spanned by $P_F \cap S$ is a collection of $O(1)$ line segments. We slide a z -vertical segment a , such that its bottom end-point always lies on $P_F \cap S$ and its top end-point slides on one of these segment on $P_C \cap T$. Consider a two-dimensional yz -parallel plane Π_a attached to segment a as above. In exactly the same manner as in Lemma 4.9, we can associate each (e, e') event that involves a with either a feature of $\mathcal{V}_1(\Gamma)$ or an event of type $G_{22,1}$, $G_{22,2}$, or $G_{22,3}$. As above, this implies the lemma. \square

4.3. The Tagansky Technique

The following lemma is proved analogously to Lemma 3.2, using the fact that the complexity of the vertical decomposition of n triangles in 3-space is $O(n^4)$ [5].

Lemma 4.12. *The number of $E_3E_3(\text{II})$ events in an arrangement of n 3-simplices in 4-space is $O(n^4)$.*

From now on, we only consider $E_3E_3(\text{I})$ events, and refer to them briefly as E_3E_3 events.

Lemma 4.13. *The number of E_3E_3 events in an arrangement of n 3-simplices in 4-space is $O(n^4\alpha(n) \log^2 n)$.*

Proof. We begin with some preliminary definitions and with the basics of the Tagansky technique.

Given two z -vertical segments e and e' that lie in a common yz -vertical plane, the *level* of the pair (e, e') is the size $|\Delta|$ of the smallest set $\Delta \subset \Gamma$, for which (e, e') is a y -vertical visibility event in $\mathcal{A}(\Gamma \setminus \Delta)$. If no such set Δ exists, the level of (e, e') is undefined. We say that (e, e') is a k -level event if it has level k .

Note that the y -vertical visibility events we have discussed so far were all 0-level events. A 1-level event (e, e') typically involves a yz -parallel trapezoid bounded by e, e', P_F , and P_C , for some $P_F, P_C \in \Gamma$, such that the interior of this trapezoid is intersected by exactly one simplex $X \in \Gamma$; if this simplex X is removed, then (e, e') becomes a (0-level) y -vertical visibility event.

Analogous definitions of levels can be formulated for all generalized events, and for features of $\mathcal{V}_1(\Gamma)$. The following lemma follows trivially from the standard Clarkson–Shor technique [4] and the bounds obtained in the preceding subsections.

Lemma 4.14. *The number of 1-level visibility events of all types other than E_3E_3 , and of all 1-level generalized events, and of all 1-level features of $\mathcal{V}_1(\Gamma)$, is $O(n^4\alpha(n) \log n)$.*

The following lemma is directly implied by the work of Tagansky [13].

Lemma 4.15. *Assume that, given any E_3E_3 event (e, e') , we can either (i) charge it to $i > 0$ 1-level E_3E_3 events, such that the number of times that each 1-level event is charged in this fashion is at most j , or (ii) charge it to a 0- or 1-level event of $\mathcal{V}_1(\Gamma)$, or a 0- or 1-level event of $\mathcal{V}_2(\Gamma)$ of any type other than E_3E_3 , or a 0- or 1-level generalized event, such that the number of times that each such event is charged in this fashion is bounded by a constant.*

- If $i/j = 2$, then the number of E_3E_3 events is $O(n^4\alpha(n) \log^2 n)$.
- If $i/j > 2$, then the number of E_3E_3 events is $O(n^4\alpha(n) \log n)$.

In the remainder of this proof we present a charging scheme for E_3E_3 events. Each such event is said to carry 4 “units of charge” (in the notation of Lemma 4.15, $i = 4$). We describe how to pass each such unit either to a 1-level E_3E_3 event or to an event of one of the types mentioned in Lemma 4.14, such that each 1-level E_3E_3 event receives at most 2 units of charge (in the notation of Lemma 4.15, $j = 2$, and hence $i/j = 2$), and each event of the other types receives $O(1)$ such units (that is, charged $O(1)$ times). Lemma 4.15 then implies that the number of E_3E_3 events is $O(n^4\alpha(n) \log^2 n)$.

We now describe the charging scheme. Consider a specific E_3E_3 event (e, e') , as depicted in Fig. 1(b,I). We slide a z -vertical segment a , such that its bottom end-point slides on $P_F \cap S \cap T$ and its top end-point slides on P_C , with the yz -parallel plane Π_a attached to it, once in each of the two possible directions (starting from $a = e$). This process is then repeated with a z -vertical segment a that slides on $P_C \cap U \cap V$ and P_F . At each of the four slidings, we carry 1 unit of charge. We specify below how this unit of charge is dispensed, when sliding in one of the directions on $P_F \cap S \cap T$ —the other three cases (one other direction on $P_F \cap S \cap T$, and two directions on $P_C \cap U \cap V$) are symmetric.

Assume, without loss of generality, that directly after the beginning of the sliding, the y -coordinate of the point $P_C \cap V$ (inside Π_a) is between the y -coordinate of $P_C \cap U$ and the y -coordinate of $P_F \cap S \cap T$. Consider the (moving) trapezoid Δ bounded by P_F, P_C, a , and the vertical line containing $P_C \cap U$. Directly after the beginning of the sliding, the interior of Δ is intersected only by V , such that V intersects the top and right edges of Δ (as in Fig. 7(e)). We pass 1 unit of charge to the first event that brings a change to

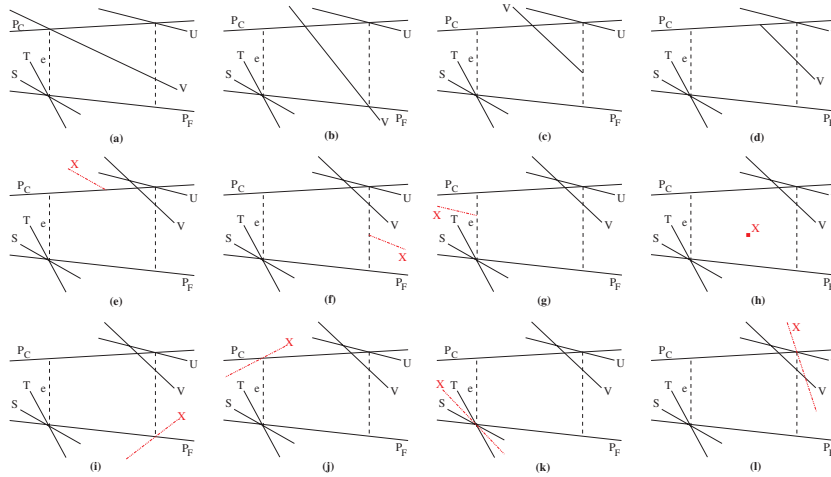


Fig. 7. Events that are charged by E_3E_3 events in arrangements of simplices during the sliding.

this situation. This can be one of the following events (all occurring within Π_a):

- (1) $P_C \cap V$ passes above $P_F \cap S \cap T$ (see Fig. 7(a)).
- (2) $V \cap P_F$ passes below $P_C \cap U$ (Fig. 7(b)).
- (3) ∂V or ∂P_F passes below $P_C \cap U$ (Fig. 7(c) shows the case of ∂V).
- (4) ∂V or ∂U intersects P_C , or ∂P_C intersects U (Fig. 7(d) shows the case of ∂V intersecting P_C).
- (5) ∂X , for some $X \in \Gamma$, intersects the top or bottom edge of Δ (Fig. 7(e) shows the case of the top edge).
- (6) ∂X , for some $X \in \Gamma$, intersects the right edge of Δ (Fig. 7(f)).
- (7) ∂X , for some $X \in \Gamma$, intersects the left edge of Δ , or ∂P_C passes above $P_F \cap S \cap T$ (Fig. 7(g) shows the former case).
- (8) A one-dimensional feature (edge) of a simplex $X \in \Gamma$ intersects the interior of Δ (Fig. 7(h)). (This happens, for example, when X first appears on the plane Π_a inside the trapezoid Δ . At the moment such an event occurs, $X \cap \Pi_a$ is a single point.)
- (9) $X \cap P_F$, for some $X \in \Gamma$, passes below $P_C \cap U$ (Fig. 7(i)).
- (10) $X \cap P_C$, for some $X \in \Gamma$, passes above $P_F \cap S \cap T$ (Fig. 7(j)).
- (11) Some $X \in \Gamma$, or ∂S , or ∂T , or ∂P_F , intersects $P_F \cap S \cap T$ (Fig. 7(k) shows the first case).
- (12) Some $X \in \Gamma$ intersects $P_C \cap U$ (Fig. 7(l)).

Claim. The number of all events of types (1)–(11) is $O(n^4 \alpha(n) \log n)$.

Proof. The claim is proved by observing that each such event, if it is the first to occur, is either a vertex of $\mathcal{A}(\Gamma)$, or a feature of $\mathcal{V}_1(\Gamma)$, or a (0- or 1-level) event of one of the types that are addressed in Lemma 4.14 or in the previous sections. An event of type (1), (7), or (10) corresponds to a feature of $\mathcal{V}_1(\Gamma)$. An event of type (2) is a $G_{3,22}$ event, while an event of type (9) is a 1-level E_3E_{22} event. An event of type (3) is a $G_{3,21}(I)$

event, while an event of type (6) is a 1-level E_3E_{21} event. An event of type (4) or (5) is a (0- or 1-level) E_3E_2 event, while an event of type (11) corresponds to a vertex of $\mathcal{A}(\Gamma)$, and an event of type (8) is a (0- or 1-level) E_3E_1 event. \square

Note that events of type (12) are 1-level E_3E_3 events, which will be charged if they are the first to occur during the sliding.

Claim. *Each 1-level E_3E_3 event receives at most 2 units of charge.*

Proof. Consider one 1-level E_3E_3 event (e, e') , formed by an interaction of $P_F \cap S \cap T$ with $P_C \cap U \cap X$, such that the trapezoid Δ_0 that has e and e' as its bases is intersected by V , for some $V \in \Gamma$ (see Fig. 7(l)). It can receive a unit of charge from four possible directions—two directions on $P_F \cap S \cap T$ and two on $P_C \cap U \cap X$. We slide a z -vertical segment a , such that its bottom end-point lies on $P_F \cap S \cap T$ and its top end-point lies on P_C , with the yz -parallel plane Π_a attached to it, starting from $a = e$ (along one of the two possible directions). Assume, without loss of generality, that directly after the beginning of the sliding, the y -coordinate of the point $P_C \cap U$ (inside Π_a) is between the y -coordinate of $P_C \cap X$ and the y -coordinate of $P_F \cap S \cap T$.

Let Δ_U be the (moving) trapezoid bounded by P_F , P_C , a , and the vertical line containing $P_C \cap U$. Directly after the beginning of the sliding, the interior of Δ_U is intersected only by V . Let (a, g) be the first event that changes the combinatorial configuration inside (or on the boundary of) Δ_U . If the top end-point of g lies on $P_C \cap U \cap V$, and the interior of Δ_U is not intersected by V (or any other simplex of Γ) when (a, g) occurs, then (a, g) is a 0-level E_3E_3 event that passes 1 unit of charge to (e, e') . In any other case, (a, g) has to be one of the events of types (1)–(12) above, and (e, e') cannot receive a unit of charge from this direction of the sliding, since the event (a, g) necessarily receives it before (figuratively speaking, (a, g) “intercepts” the unit of charge that can come from this direction).

This means that (e, e') receives a unit of charge from this specific direction of sliding if and only if the first event that changes the combinatorial configuration inside (and on the boundary of) Δ_U is the event of V intersecting $P_C \cap U$. An analogous argument can be repeated for the other three possible directions of sliding from (e, e') —one on $P_F \cap S \cap T$ and two on $P_C \cap U \cap X$.

The event (e, e') can be of four types:

- A. ∂V lies inside Δ_0 (see Fig. 8(a)).
- B. V intersects neither e nor e' (Fig. 8(b)).
- C. V intersects both e and e' (Fig. 8(c)).
- D. V intersects e , but not e' (Fig. 8(d)); this case is symmetric to the case in which V intersects e' but not e .

We analyze each of these types of 1-level events in turn, and prove that events of types A, B, and C cannot be charged at all, while events of type D receive at most 2 units of charge.

Suppose the event (e, e') is of type A. We slide a z -vertical segment a , such that its bottom end-point lies on $P_F \cap S \cap T$ and its top end-point lies on P_C , starting from $a = e$, as above. To show that (e, e') cannot receive a unit of charge from this direc-




Fig. 8. The types of 1-level E_3E_3 events.

tion of sliding, it is sufficient to show that the first event that changes the combinatorial configuration inside (or on the boundary of) Δ_U cannot be the event of V intersecting $P_C \cap U$ such that V does not intersect the interior of Δ_U . This is indeed easy to see, since before V can intersect $P_C \cap U$ in this fashion, ∂V has to cease lying inside Δ_U (assuming general position). Thus, before V intersects $P_C \cap U$ as above, there occurs an event of ∂V exiting the interior of Δ_U , which changes the configuration inside Δ_U . This means that the event of V intersecting $P_C \cap U$ as above cannot be the first event that is encountered. Since this argument symmetrically applies to all four directions of the sliding, it implies that no event of type A can be charged at all.

Suppose the event (e, e') is of type B or C. We can assume that the first event (a, g) that changes the combinatorial configuration inside (or on the boundary of) Δ_U is the event of V intersecting $P_C \cap U$ such that V does not intersect the interior of Δ_U (because otherwise (e, e') cannot be charged from this direction, as explained above). At the beginning of the sliding, V separates two vertices of Δ_U from two other vertices of Δ_U . By the assumption that was just made, the first vertex of Δ_U intersected by V is $P_C \cap U$. This means that at the moment the event (a, g) materializes, V separates one vertex of Δ_U from two other vertices of Δ_U . Thus, a part of V necessarily lies inside Δ_U at the time of this event, which contradicts the assumption, and implies as above that (a, g) cannot pass a unit of charge to (e, e') at all. Since this argument symmetrically applies to all four directions of the sliding, it proves that no event of type B or C can be charged at all.

Suppose the event (e, e') is of type D, and that V intersects e , but not e' (the analysis for the case in which V intersects e' but not e is fully symmetric). Assume that the first event (a, g) that is encountered is as above. At the beginning of the sliding, V intersects a , and the assumption that was just made implies that this is still true at the time the event (a, g) materializes. This contradicts the assumption and implies as above that (e, e') cannot be charged by sliding on $P_F \cap S \cap T$ (in any of the two directions, since the argument symmetrically applies for the other direction of sliding as well). This does not rule out the possibility that (e, e') is charged by sliding on $P_C \cap U \cap X$. Thus, events of type D can receive at most 2 units of charge (from the two directions of sliding on $P_C \cap U \cap X$). \square

As already explained, combining this with Lemma 4.15 completes the proof of Lemma 4.13. \square

We have now bounded by $O(n^4 \alpha(n) \log^2 n)$ the number of visibility events of all possible combinatorial types. This concludes the proof of Theorem 4.1. \square

Acknowledgements

I thank Micha Sharir for his invaluable help in preparing this paper, as well as Danny Halperin for introducing me to vertical decompositions in four dimensions and for subsequent discussions.

References

1. 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, North-Holland, Amsterdam, 2000.
2. B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties and its applications. *Theoretical Computer Science*, 84:77–105, 1991. Also in *Proc. International Colloquium on Automata, Languages and Programming*, pages 179–193, 1989.
3. K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM Journal on Computing*, 17:830–847, 1988.
4. K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete & Computational Geometry*, 4:387–421, 1989.
5. M. de Berg, L. J. Guibas, and D. Halperin. Vertical decompositions for triangles in 3-space. *Discrete & Computational Geometry*, 15:35–61, 1996.
6. H. Edelsbrunner. The upper envelope of piecewise linear functions: tight complexity bounds in higher dimensions. *Discrete & Computational Geometry*, 4:337–343, 1989.
7. H. Edelsbrunner, R. Seidel, and M. Sharir. On the zone theorem for hyperplane arrangements. *SIAM Journal on Computing*, 22(2):418–429, 1993.
8. L. J. Guibas, D. Halperin, J. Matoušek, and M. Sharir. On vertical decomposition of arrangements of hyperplanes in four dimensions. *Discrete & Computational Geometry*, 14:113–122, 1995.
9. D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 21, pages 389–412. CRC Press, Boca Raton, FL, 1997.
10. V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *Journal of the ACM*, to appear.
11. V. Koltun and C. Wenk. Matching polyhedral terrains using overlays of envelopes. *Algorithmica*, to appear.
12. M. Sharir and P. K. Agarwal. *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
13. B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete & Computational Geometry*, 16:455–479, 1996.

Received April 26, 2002, and in revised form May 27, 2003. Online publication February 18, 2004.