

A New Technique for Analyzing Substructures in Arrangements of Piecewise Linear Surfaces*

B. Tagansky

School of Mathematical Sciences, Tel Aviv University,
Tel Aviv 69978, Israel
mia@math.tau.ac.il

Abstract. We present a simple but powerful new probabilistic technique for analyzing the combinatorial complexity of various substructures in arrangements of piecewise-linear surfaces in higher dimensions. We apply the technique (a) to derive new and simpler proofs of the known bounds on the complexity of the lower envelope, of a single cell, or of a zone in arrangements of simplices in higher dimensions, and (b) to obtain improved bounds on the complexity of the vertical decomposition of a single cell in an arrangement of triangles in 3-space, and of several other substructures in such an arrangement (the entire arrangement, all nonconvex cells, and any collection of cells). The latter results also lead to improved algorithms for computing substructures in arrangements of triangles and for translational motion planning in three dimensions.

1. Introduction

The study of arrangements of curves or surfaces is an important area of research in computational and combinatorial geometry, because many geometric problems in diverse areas can be reduced to problems involving arrangements. We assume in this paper the reader is familiar with basic terminology and results concerning arrangements, and refer to [14], [18], and [20] for more details.

In this paper we present a new technique for analyzing the combinatorial complexity of various substructures in arrangements of piecewise-linear surfaces in higher dimensions. The technique uses a simplified variant of the Clarkson–Shor probabilistic analysis technique [11], and combines it with local geometric analysis of the arrangement under

* Work on this paper was partially supported by a grant from the G.I.F., the German–Israeli Foundation for Scientific Research and Development. The research reported in this paper is part of the author’s Ph.D. thesis prepared under the supervision of Prof. Micha Sharir.

consideration. The technique is a simplification of a previous technique that was used for analyzing the complexity of lower envelopes and other substructures in arrangements of more general algebraic surfaces (see [2], [19], and [24]). Our technique also somewhat resembles the inductive technique used in [3], [5], and [16] for analyzing arrangements of piecewise-linear surfaces, but is more general and more powerful.

We first exemplify the technique in Section 2 on the problem of bounding the maximum complexity of the lower envelope of n line segments in the plane. We obtain a bound $O(n \log n)$, which is slightly weaker than the tight bound of $\Theta(n\alpha(n))$ [21, [26]. The proof however is very short and simple. (Intuitively, the new technique is not “sensitive” enough to yield bounds with factors like $\alpha(n)$.)

We then describe, in Section 3, the technique in a general abstract setting, and then apply this general methodology to several specific problems.

The first application, given in Sections 4 and 5, is to arrangements of n $(d - 1)$ -simplices in \mathbb{R}^d . We obtain new (and simpler) proofs of the known bounds for the maximum complexity of the lower envelope, of a single cell, and of a zone in such an arrangement of simplices. For lower envelopes, the bound is $\Theta(n^{d-1}\alpha(n))$; previous proofs of this bound are given in [15], [23], and [25]. For a single cell or a zone of a fixed-degree algebraic surface or an arbitrary convex surface, the bound is $O(n^{d-1} \log n)$; a previous, and considerably more involved, proof is given in [5].

We then derive, in Section 6, an improved bound on the complexity of vertical decompositions in arrangements of n triangles in 3-space. The previous bound in [13] for the complexity of the vertical decomposition of a single cell in such an arrangement was $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε . We improve this bound to $O(n^2 \log^2 n)$ (actually, this bound applies also to the vertical decomposition of the zone of any fixed-degree algebraic surface or the surface of any convex body). We then extend the result to several more complex portions of the arrangement: We show that the complexity of the vertical decomposition of the entire arrangement is $O(K + n^2\alpha(n) \log n)$, where K is the complexity of the undecomposed arrangement; the previous bound of [13] was $O(K + n^{2+\varepsilon})$, for any $\varepsilon > 0$. We also obtain a bound of $O(n^{2.5})$ for the complexity of the vertical decomposition of all nonconvex cells of the arrangement, and a bound of $O(m^{1/3}n^2 + n \log^2 n)$ for the complexity of the vertical decomposition of any m cells of the arrangement. We are not aware of any previous subcubic bounds for these two latter complexities. Still, we do not know whether the last two bounds are close to optimal; the actual complexities of the (undecomposed) cells in question are known to be smaller.

Using the randomized algorithm of [12] and applying our new bounds in its analysis, we obtain improved bounds for constructing the relevant substructures. For example, (the vertical decomposition of) a single cell in an arrangement of n triangles in 3-space can be computed in $O(n^2 \log^3 n)$ expected time. Also, we obtain an algorithm for constructing all nonconvex cells of the arrangement, whose running time is close to $O(n^{2.5})$. This yields a simple motion-planning algorithm, of the same complexity, for an arbitrary polyhedron translating in a three-dimensional polyhedral environment.

This paper does not exhaust all currently known applications of our new technique. In several companion papers [8], [9] we apply the technique to obtain tight complexity bounds for the union of axis-parallel hypercubes in higher dimensions, and tight or nearly tight bounds for various generalized Voronoi diagrams in higher dimensions.

Other applications of our technique improved bounds on the complexity of the union of convex polyhedra in 3-space, improving by a logarithmic factor earlier bounds given in [6] and [7].

2. Lower Envelope of Segments

We begin by exemplifying our technique on the problem of bounding the maximum complexity of the lower envelope of n line segments in the plane.

Let \mathcal{L} be a set of n nonvertical line segments in the plane, in general position. By general position we mean that any two segments intersect in at most a single point, and that no three segments share a single point. Let $\mathcal{A}(\mathcal{L})$ denote the arrangement of \mathcal{L} . An intersection point p between the interiors of two segments of \mathcal{L} is said to be a k -level inner vertex of $\mathcal{A}(\mathcal{L})$ if there are exactly k segments in \mathcal{L} passing strictly below p . (We refer to such vertices as *inner vertices*, to distinguish them from the endpoints of the segments.) Let $C_k(\mathcal{L})$ denote the number of k -level inner vertices in $\mathcal{A}(\mathcal{L})$, and let $C_k(n)$ denote the maximum of $C_k(\mathcal{L})$ over all sets of n segments in the plane, in general position. We want to bound $C_0(n)$, the maximum number of inner vertices on the lower envelope of \mathcal{L} . It is well known that $C_0(n) = \Theta(n\alpha(n))$ [21], [26]; here we only prove a weaker upper bound of $O(n \log n)$. This serves as an introductory example of our technique for analyzing the complexity of various substructures in arrangements.

Our technique consists of two parts. The first part involves geometric analysis that enables us to bound the number of 0-level inner vertices in terms of the number of 1-level inner vertices, plus some excess that we can bound using a more direct analysis (which is trivial in the case of segments). It proceeds as follows.

Let p be a 0-level inner vertex of $\mathcal{A}(\mathcal{L})$. We sweep a vertical line from p to the right, and stop the sweeping as soon as we encounter one of the following two types of events (see Fig. 1):

- (i) We reach an endpoint of one of the segments.
- (ii) We reach a 1-level vertex.

An important observation is that the sweeping line can never reach a 0-level inner vertex before any of these events are encountered, so we can reach every such event only from (at most) one 0-level vertex. Let $D(\mathcal{L})$ denote the number of events of type (i); we clearly have $D(\mathcal{L}) \leq 2n$, which leads to the inequality

$$C_0(\mathcal{L}) \leq C_1(\mathcal{L}) + D(\mathcal{L}) \leq C_1(\mathcal{L}) + 2n. \quad (1)$$

The second part of our technique applies a simple variant of the random-sampling technique of Clarkson and Shor [11]. Together with inequality (1), this gives a recurrence relation whose solution yields a bound on $C_0(n)$.

This part proceeds as follows. We take a random sample \mathcal{R} of $n - 1$ segments of \mathcal{L} , that is, we obtain \mathcal{R} by removing a random segment from \mathcal{L} . Each 0-level inner vertex p of $\mathcal{A}(\mathcal{L})$ will appear as a 0-level vertex of $\mathcal{A}(\mathcal{R})$ if and only if the two intersecting segments containing p are in \mathcal{R} . This will happen with probability $(n - 2)/n$. Similarly,




Fig. 1. Sweeping from a 0-level inner vertex: terminal events of types (i) and (ii).

each 1-level inner vertex p' of $\mathcal{A}(\mathcal{L})$ will appear as a 0-level vertex of $\mathcal{A}(\mathcal{R})$ if and only if the segment in \mathcal{L} below p' is not chosen in \mathcal{R} . This will happen with probability $1/n$. No other inner vertex of $\mathcal{A}(\mathcal{L})$ can appear as a 0-level vertex of $\mathcal{A}(\mathcal{R})$. Hence

$$\mathbf{E}(C_0(\mathcal{R})) = \frac{n-2}{n}C_0(\mathcal{L}) + \frac{1}{n}C_1(\mathcal{L}), \tag{2}$$

where \mathbf{E} denotes expectation with respect to the random sample \mathcal{R} . Putting (1) and (2) together, we get

$$\begin{aligned} \frac{n-1}{n}C_0(\mathcal{L}) &= \frac{n-2}{n}C_0(\mathcal{L}) + \frac{1}{n}C_0(\mathcal{L}) \\ &\leq \frac{n-2}{n}C_0(\mathcal{L}) + \frac{1}{n}(C_1(\mathcal{L}) + 2n) \\ &= \mathbf{E}(C_0(\mathcal{R})) + 2. \end{aligned}$$

Since $\mathbf{E}(C_0(\mathcal{R})) \leq C_0(n-1)$, we obtain the following recurrence:

$$\frac{n-1}{n}C_0(n) \leq C_0(n-1) + 2, \tag{3}$$

and $C_0(2) = 1$. The solution of this recurrence is easily seen to be $C_0(n) \leq 1 + 2n \log_e(n-1)$, so we obtain:

Theorem 2.1. *The number of inner vertices on the lower envelope of n segments in the plane is at most $2n \log_e n$.*

Remark. The same bound for $C_0(n)$ has been obtained by a different proof, given in [25] and based on analysis of Davenport–Schinzel sequences of order 3. Both proofs obtain essentially the same recurrence for $C_0(n)$. An intriguing open problem is to obtain a purely geometric (and hopefully simple) proof of the known tight bound $\Theta(n \alpha(n))$.

3. The Technique in an Abstract Setting

In this section we describe our analysis technique in a more general and abstract manner, so as to facilitate its subsequent application to several different problems. It would be helpful to compare the general setting with the specific example given above.

The general setup is similar to that used by Clarkson [10], Clarkson and Shor [11], de Berg *et al.* [12], Mulmuley [22], and others. Let \mathcal{L} be a set of n (geometric) objects. Let $\Phi(\mathcal{L})$ be a set of “features,” where each feature is defined in terms of d objects of \mathcal{L} . We refer to d as the *combinatorial dimension* of the problem. In addition to the d objects defining a feature φ , there may be other objects that “conflict” with φ . If no object conflicts with φ , then φ is said to be *exposed*, or to lie at level 0. Otherwise, we define the *level* of φ to be the smallest number of objects whose removal causes φ to become exposed (observe that none of the removed objects can be any of the d objects defining φ). Let $\Phi_k(\mathcal{L})$ denote the subset of features in $\Phi(\mathcal{L})$ whose level is exactly k . An important characteristic property that we use is that for any $\mathcal{R} \subset \mathcal{L}$ we have $\Phi(\mathcal{R}) \cap \Phi_0(\mathcal{L}) \subseteq \Phi_0(\mathcal{R})$. This states the obvious property that, if a feature $\varphi \in \Phi(\mathcal{L})$ is defined only by objects in \mathcal{R} and does not conflict with any object of \mathcal{L} , then it does not conflict with any object of \mathcal{R} and is thus in $\Phi_0(\mathcal{R})$. Our goal is to bound $|\Phi_0(\mathcal{L})|$ as a function of n .

Remark. In most applications the minimal set of objects whose removal exposes a feature φ is unique, and consists of all objects that conflict (individually) with φ . In some cases, though, exposing φ can be done in more than one way. As an example, consider the case where \mathcal{L} is again a set of segments in the plane, and where a vertex of the arrangement of \mathcal{L} is exposed if it lies on the boundary of the face of the arrangement that contains the origin. In this case it may be possible to expose a vertex by removing several different subsets of segments. (See also [12]).

Our analysis technique consists of two parts. The first part, which is specific for the problem under consideration, uses geometric analysis to bound the number of 0-level features in terms of the number of 1-level features plus some known excess. The result of this part is an inequality of the form

$$\rho C_0(\mathcal{L}) \leq C_1(\mathcal{L}) + D(\mathcal{L}), \tag{4}$$

for some constant $\rho > 0$, where $C_k(\mathcal{L})$ is the number of k -level features in $\Phi(\mathcal{L})$, and $D(\mathcal{L})$ is the excess. We assume that $D(\mathcal{L})$ can be bounded, using separate analysis, by some function $D(n)$ of n .

The second part of our technique is based, as in the case of segments, on a simple variant of the random-sampling technique of Clarkson and Shor [11]. Together with the above inequality, this gives us a recurrence relationship for $C_0(n)$, the maximum of $C_0(\mathcal{L})$ over all collections \mathcal{L} of n objects of the type that we consider here. The solution of this recurrence yields the desired bound.

As in the case of segments, we take a random sample $\mathcal{R} \subseteq \mathcal{L}$ of size $n - 1$, that is, we obtain \mathcal{R} by removing a random object from \mathcal{L} . Each 0-level feature φ of $\Phi(\mathcal{L})$ will appear as a 0-level feature of $\Phi(\mathcal{R})$ if (and only if) the d objects defining φ are chosen in \mathcal{R} . This will happen with probability $(n - d)/n$. For each 1-level feature φ' of $\Phi(\mathcal{L})$, there exists at least one conflicting object whose removal makes φ' a 0-level feature. Hence the probability that φ' will appear as a 0-level feature of $\Phi(\mathcal{R})$ is at least $1/n$. Hence we have

$$\mathbf{E}(C_0(\mathcal{R})) \geq \frac{n - d}{n} C_0(\mathcal{L}) + \frac{1}{n} C_1(\mathcal{L}), \tag{5}$$

where \mathbf{E} denotes expectation with respect to the random sample \mathcal{R} . Putting (4) and (5) together, we obtain

$$\begin{aligned} \frac{n-d+\rho}{n}C_0(\mathcal{L}) &= \frac{n-d}{n}C_0(\mathcal{L}) + \frac{\rho}{n}C_0(\mathcal{L}) \\ &\leq \frac{n-d}{n}C_0(\mathcal{L}) + \frac{1}{n}(C_1(\mathcal{L}) + D(\mathcal{L})) \\ &\leq \mathbf{E}(C_0(\mathcal{R})) + \frac{1}{n}D(\mathcal{L}). \end{aligned}$$

Since $\mathbf{E}(C_0(\mathcal{R})) \leq C_0(n-1)$, we get the following recurrence:

$$\frac{n-d+\rho}{n}C_0(n) \leq C_0(n-1) + \frac{1}{n}D(n). \quad (6)$$

This recurrence may have different solutions, depending on the possible values of ρ and $D(n)$. In general, we have:

Proposition 3.1. *The solution $F(n)$ of the recurrence*

$$\frac{n-t}{n}F(n) \leq F(n-1) + \frac{1}{n}D(n), \quad (7)$$

for any fixed real $t \geq 0$, satisfies

$$F(n) = O\left(n^t \cdot \left[1 + \sum_{j=\lceil t \rceil+1}^n \frac{D(j)}{j^{t+1}}\right]\right). \quad (8)$$

In particular,

- (a) If $D(n) = O(n^{t'})$, for any $t' < t$, then $F(n) = O(n^t)$.
- (b) If $D(n) = O(n^t \log^c n)$, for any $c \geq 0$, then $F(n) = O(n^t \log^{c+1} n)$.
- (c) If $D(n) = O(n^{t'})$, for any $t' > t$, then $F(n) = O(n^{t'})$.

Proof. Claims (a)–(c) follow from (8) by routine calculations, which we omit here. We thus only give the proof of (8). The proof is trivial for integer t : Dividing (7) by $(n-1)(n-2)\cdots(n-t)$, we obtain a telescopic recurrence whose solution is easily seen to satisfy (8). The proof is more involved when t is not an integer, and proceeds as follows.

For any positive real number t , and for any $n > t$, define the factorial-like function

$$f_t(n) = n(n-1)\cdots(n-\lfloor t \rfloor + 1)(n-\lfloor t \rfloor)^{t-\lfloor t \rfloor} = \Theta(n^t),$$

and let $f_0(n) = 1$. First we show that for all positive integers $k < n-t$ we have

$$\frac{f_t(n)}{f_t(n-k)} \geq \frac{f_t(n)}{f_t(n-t)}.$$

The proof proceeds by induction on k . Let $m = \lfloor t \rfloor$ and $r = t - m$. For $k = 1$ we have

$$\begin{aligned} \frac{f_t(n)}{f_t(n-1)} &= \frac{n}{n-m} \left(\frac{n-m}{n-m-1} \right)^r = \frac{n}{n-m} \left(1 - \frac{1}{n-m} \right)^{-r} \\ &\geq \frac{n}{n-m} \left(1 - \frac{r}{n-m} \right)^{-1} = \frac{n}{n-t} = \frac{f_1(n)}{f_1(n-t)} \end{aligned}$$

(the inequality holds since $0 \leq r < 1$). For $k \geq 2$, we have, by the induction hypothesis (applied to $k - 1$ and to 1),

$$\begin{aligned} \frac{f_t(n)}{f_t(n-k)} &= \frac{f_t(n)}{f_t(n-k+1)} \cdot \frac{f_t(n-k+1)}{f_t(n-k)} \\ &\geq \frac{f_{k-1}(n)}{f_{k-1}(n-t)} \cdot \frac{n-k+1}{n-k+1-t} = \frac{f_k(n)}{f_k(n-t)}, \end{aligned}$$

as long as $k < n - t$. Next we unfold the recurrence formula, to obtain,

$$\begin{aligned} F(n) &\leq \frac{n}{n-t} \left(F(n-1) + \frac{D(n)}{n} \right) \\ &\leq \frac{n}{n-t} \left[\frac{n-1}{n-t-1} \left(F(n-2) + \frac{D(n-1)}{n-1} \right) + \frac{D(n)}{n} \right] \\ &\quad \vdots \\ &\leq \frac{f_k(n)}{f_k(n-t)} \cdot F(n-k) + \sum_{i=1}^k \frac{f_i(n)}{f_i(n-t)} \cdot \frac{D(n-i+1)}{n-i+1} \\ &\leq \frac{f_t(n)}{f_t(n-k)} \cdot F(n-k) + \sum_{i=1}^k \frac{f_i(n)}{f_i(n-i)} \cdot \frac{D(n-i+1)}{n-i+1}. \end{aligned}$$

Taking $k = n - \lceil t \rceil$ (which satisfies the constraint $k < n - t$), we get

$$\frac{F(n)}{f_t(n)} \leq \frac{F(\lceil t \rceil)}{f_t(\lceil t \rceil)} + \sum_{j=\lceil t \rceil+1}^n \frac{D(j)}{j f_t(j-1)}.$$

Since $f_t(j-1) = \Theta(j^t)$, the bound of (8) follows. This completes the proof of the proposition. \square

4. Lower Envelopes in Arrangements of Simplices

Let \mathcal{S} be a set of $n(d-1)$ -simplices in \mathbb{R}^d , and let $\mathcal{A}(\mathcal{S})$ denote the arrangement of \mathcal{S} . In this section we consider the problem of bounding the combinatorial complexity of the lower envelope of \mathcal{S} . That is, if we regard each simplex in \mathcal{S} as a partial linear $(d-1)$ -variate function, the lower envelope is the (graph of the) pointwise minimum of these functions. The lower envelope is a polyhedral set, and we measure its complexity by the

number of vertices. (The number of all other faces is easily seen to be proportional to the number of vertices, if the simplices are in general position. We do indeed assume that the simplices are in general position, arguing, as in previous works, that the maximum complexity of the envelope is attained for sets \mathcal{S} in general position.) Note that the pointwise minimum of simplices is generally not continuous.

We derive here a new, and simpler proof of the known tight bound on the complexity of the envelop, previously proved in [15], [23], and [25]:

Theorem 4.1. *The maximum number of vertices of the lower envelope of n $(d - 1)$ -simplices in \mathbb{R}^d , for $d \geq 2$, is $\Theta(n^{d-1}\alpha(n))$.*

Proof. We draw a vertical hyperplane through each $(d - 2)$ -face of every simplex in \mathcal{S} , and add those hyperplanes to our arrangement. An *outer* vertex in the augmented arrangement is a vertex that lies in one of these vertical hyperplanes, whereas an *inner* vertex lies in the intersection of exactly d simplex interiors and avoids these hyperplanes.

We assume that the simplices of \mathcal{S} are bounded and in *general position*. It is easily seen that neither assumption involves any real loss of generality (see also the previous cited proofs).

The lower bound has been shown, e.g., in [23]; it follows easily from the lower bound $\Omega(n\alpha(n))$ for arrangements of segments in the plane, as given in [26].

We prove the upper bound by induction on d . The bound holds for $d = 2$ (the case of line segments in the plane), as follows from the results of [21]. Fix $d \geq 3$, assume that the upper bound holds for all $2 \leq d' \leq d - 1$, and let \mathcal{S} be a collection of n $(d - 1)$ -simplices in \mathbb{R}^d .

A vertex $v \in \mathcal{A}(\mathcal{S})$ is said to be a *k-level* vertex if there are exactly k simplices of \mathcal{S} passing strictly below v . Let $C_k(\mathcal{S})$ denote the number of k -level inner vertices in $\mathcal{A}(\mathcal{S})$, and let $C_k(n, d)$ denote the maximum of $C_k(\mathcal{S})$ over all possible sets \mathcal{S} of n $(d - 1)$ -simplices in \mathbb{R}^d in general position. Let $D(\mathcal{S})$ denote the number of outer vertices of the lower envelope, and let $D(n, d)$ denote the maximum of $D(\mathcal{S})$ over all sets \mathcal{S} of n simplices as above. Our goal is thus to bound $C_0(n, d)$ and $D(n, d)$, which together serve as our measure of the combinatorial complexity of the lower envelope.

The outer vertices are easy to deal with. Let f be a $(d - 2)$ -face of a simplex s , and let h_f be the vertical hyperplane passing through f . We intersect h_f with all the simplices in \mathcal{S} , thereby obtaining a $(d - 1)$ -dimensional arrangement of $O(n)$ $(d - 2)$ -simplices within h_f , which we denote by \mathcal{A}_f (the intersection of a simplex with a hyperplane is not necessarily a simplex, but it is a convex polytope of constant complexity, and can be partitioned into $O(1)$ simplices, where this constant depends on d). We also denote by \mathcal{A}_f^- the arrangement \mathcal{A}_f after removing the intersection $s \cap h_f$ from it. Then each outer vertex of the lower envelope of $\mathcal{A}(\mathcal{S})$ within h_f can be identified with an (inner or outer) vertex of the lower envelope of either \mathcal{A}_f or \mathcal{A}_f^- . Since there are $nd(d - 1)/2$ such $(d - 2)$ -faces f , we get

$$D(n, d) \leq \frac{nd(d - 1)}{2} \cdot 2(C_0(O(n), d - 1) + D(O(n), d - 1)),$$

and the induction hypothesis implies that $D(n, d) = O(n^{d-1}\alpha(n))$.

Next we bound the number of inner vertices, using our new technique. Each 0-level inner vertex v of $\mathcal{A}(\mathcal{S})$ is incident to exactly d edges of $\mathcal{A}(\mathcal{S})$, each of which emanates from v and is hidden from below by one of the simplices incident to v . We trace each of these edges from v , and stop tracing as soon as we first encounter one of the following types of events:

- (i) We reach a vertical $(d - 1)$ -dimensional slab h , spanned by a $(d - 2)$ -face of some simplex in \mathcal{S} .
- (ii) We meet a new simplex at a point v' , which is necessarily a 1-level inner vertex of $\mathcal{A}(\mathcal{S})$. In this case we say that v' and v are neighbors (in the arrangement $\mathcal{A}(\mathcal{S})$).

An important observation is that we can never reach a 0-level vertex during any of these d tracings, before one of these types of events is encountered. Hence any of these “terminal” events can be reached at most once along any of its incident edges.

If we intersect the $nd(d - 1)/2$ vertical slabs that arises in case (i) above with the simplices in \mathcal{S} , we get $nd(d - 1)/2$ arrangements, each consisting of $O(n)$ $(d - 2)$ -simplices in \mathbb{R}^{d-1} . Each of the terminal events of type (i) is easily seen to be either a 1-level or a 2-level inner vertex of one of these arrangements. We charge each such vertex by, say 2 units from each of the two possible directions in which it can be reached, for a total of 4 units. The total charge of all terminal events of types (i) and (ii) is thus bounded by

$$\frac{nd(d - 1)}{2} \cdot 4(C_1(O(n), d - 1) + C_2(O(n), d - 1)).$$

Using a standard application of the Clarkson–Shor probabilistic technique [11] (with a random sample of, say, half of the simplices in such a $(d - 1)$ -dimensional arrangement), we easily obtain

$$C_1(O(n), d - 1) + C_2(O(n), d - 1) = O(C_0(O(n), d - 1)).$$

Thus, by the induction hypothesis, the total charge of all terminal events of type (i) is bounded by $O(n^{d-1}\alpha(n))$.

In case of a type (ii) terminal event, we charge the 1-level inner vertex v' that we reach. Let $m(v')$ denote the number of 0-level neighbors of v' , which is the number of times v' will be charged. We give v' a total of 1 unit of charge, and it pays $1/m(v')$ units to each 0-level neighbor. A 1-level inner vertex can be charged by type (ii) events up to d times, thus the charge it can pay to a 0-level neighbor will be at least $1/d$ units. This implies that any 0-level inner vertex v receives at least 1 unit of charge.

We next claim that every 0-level inner vertex v receives at least $\frac{4}{3}$ units of charge. By construction, this is the case when at least one of the tracings from v reaches a terminal event of type (i), so we only consider vertices v for which all d tracings from v terminate at type (ii) events.

We represent each inner vertex of $\mathcal{A}(\mathcal{S})$ by the (unordered) tuple of the d simplices incident to the vertex. Let v be a 0-level inner vertex represented by the tuple $(1, 2, \dots, d)$.

For any simplex x , not incident to v , let V_x be the collection of all 1-level neighbors of v incident to x , let S_x be the set of simplices hiding some vertex in V_x , and let N_x

be the set of 0-level neighbors of the vertices in V_x (note that $|S_x| = |V_x|$). We have the following two claims:

- (a) Any member of $N_x \setminus \{v\}$ is represented by $(1, \dots, j-1, x, j+1, \dots, d)$ for some $j \in \{1, \dots, d\} \setminus S_x$.

Indeed, any vertex $v' \in V_x$ is obtained by replacing one of the simplices incident to v , say simplex 1, by x , where the replaced simplex 1 hides v' from below. We represent such a vertex v' by the tuple $(x, 2, \dots, d|1)$. To prove the claim, let $v' \in V_x$ have this representation, and suppose that v' has another 0-level neighbor v_1 . The representation of any such 0-level neighbor is obtained by removing one of the simplices $2, \dots, d$ incident to v' and by replacing it with the hiding simplex 1. Suppose that v_1 is not incident to the simplex 2; then v_1 is represented by the tuple $(x, 1, 3, \dots, d)$. Now v cannot have a 1-level neighbor incident to x and hidden by simplex 2, since the representation of such a neighbor would have to be the same as that of v_1 , which is not a 1-level vertex. Thus the simplex 2 cannot be in S_x . The claim now follows.

- (b) For any $k \in \{1, \dots, d\}$, only one set N_x can have a member that is not incident to the simplex k .

To see this, consider, with no loss of generality, the case $k = 1$, and assume to the contrary that $v'_1 \in V_x$ has a 0-level neighbor v_1 represented by $(x, 2, \dots, d)$, and that $v'_2 \in V_y$ has a 0-level neighbor v_2 represented by $(y, 2, \dots, d)$, and $x \neq y$. (It follows from claim (a) that these are the only possible vertices with these properties.) The three vertices v, v_1, v_2 lie on the same line segment e , formed by the intersection of the simplices $2, \dots, d$, and none of them is hidden from below by any simplex of \mathcal{S} . Suppose, with no loss of generality, that v_1 is the vertex in the middle. Then the hyperplane π_x containing the simplex x must pass strictly below one of the vertices v, v_2 ; suppose, for definiteness, that it passes below v_2 . Consider the polygonal path $p: v_1 v'_1 v v'_2 v_2$. By construction, p starts on simplex x , and it does not meet any vertical slab through any $(d - 2)$ -face of x . Thus p must reach v_2 while it is above the simplex x , a contradiction that establishes our claim. See Fig. 2 for an illustration. Similar arguments apply when v or v_2 is the middle vertex.

Suppose that the 1-level neighbors of v are grouped into m subsets V_{x_1}, \dots, V_{x_m} , for appropriate distinct simplices x_1, \dots, x_m . We clearly have $\sum_{i=1}^m |V_{x_i}| = d$. Claim (a) implies that $|V_{x_i}| + |N_{x_i} \setminus \{v\}| \leq d$, for each i , and Claim (b) implies that $\sum_{i=1}^m |N_{x_i} \setminus \{v\}| \leq d$. The number of 0-level neighbors of each member of V_{x_i} is at most $|N_{x_i}|$, so each of these members will pay to v at least $1/|N_{x_i}|$ units. Summing this over all 1-level neighbors




Fig. 2. The polygonal path used in the proof of Claim (b); it is illustrated in \mathbb{R}^3 , with a side view (left) and a top view (right).

of v , we conclude that the total charge that v will receive is at least $\sum_{i=1}^m |V_{x_i}|/|N_{x_i}|$. We thus want to minimize the sum $\sum_{i=1}^m a_i/(b_i + 1)$, subject to the constraints (i) a_i are positive integers, (ii) b_i are nonnegative integers, (iii) $\sum_{i=1}^m a_i = d$, (iv) $a_i + b_i \leq d$, for each $i = 1, \dots, m$, and (v) $\sum_{i=1}^m b_i \leq d$. A simple (but formally somewhat tedious) exercise shows that the total charge is minimized when $m = 2$, $a_1 = b_2 = \lfloor d/2 \rfloor$, and $a_2 = b_1 = \lceil d/2 \rceil$. In this case v receives

$$\left\lfloor \frac{d}{2} \right\rfloor \cdot \frac{1}{\lfloor d/2 \rfloor + 1} + \left\lceil \frac{d}{2} \right\rceil \cdot \frac{1}{\lceil d/2 \rceil + 1}$$

units of charge. For $d = 2r$ or $d = 2r - 1$, this amount is $1 + (r - 1)/(r + 1)$, which is at least $\frac{4}{3}$ for any $d \geq 3$. To summarize, we have shown that each 0-level inner vertex receives at least $\frac{4}{3}$ units of charge. This implies that

$$\frac{4}{3}C_0(\mathcal{S}) \leq C_1(\mathcal{S}) + O(n^{d-1}\alpha(n)).$$

Using our technique, we obtain the recurrence

$$\frac{n - d + \frac{4}{3}}{n}C_0(n) \leq C_0(n - 1) + \frac{1}{n}O(n^{d-1}\alpha(n)),$$

whose solution is $C_0(n) = O(n^{d-1}\alpha(n))$ (see Proposition 3.1). This completes the proof of the theorem. □

Remark. This analysis extends the one give in Section 2 for the case of segments in the plane. It is interesting to note that the argument given above fails for $d = 2$; that is, we can only show in that case that a 0-level vertex gets 1 unit of charge, which leads to the logarithmic factor in the bound. This indicates that the planar case is much subtler than the higher-dimensional cases, and that our technique is not sensitive enough to yield the bound $O(n\alpha(n))$ for the planar case.

5. Single Cell or Zone in an Arrangement of Simplices

Let \mathcal{S} be a set of n $(d - 1)$ -simplices in \mathbb{R}^d . As in the preceding section, we assume that the simplices of \mathcal{S} are in general position. As argued in [4] and [5], this involves no loss of generality in the foregoing analysis.

Let P be a fixed algebraic surface of dimension $\leq d - 1$ and of constant algebraic degree, or the boundary of an arbitrary convex set. We denote by $Z_P(\mathcal{S})$ the *zone* of P in $\mathcal{A}(\mathcal{S})$, namely the collection of all cells of $\mathcal{A}(\mathcal{S})$ whose intersection with P is nonempty. We want to bound the combinatorial complexity of $Z_P(\mathcal{S})$, which we measure by the number of vertices of the boundary of $Z_P(\mathcal{S})$ (the number of all other faces of the boundary is clearly proportional to the number of vertices, assuming general position of the simplices of \mathcal{S} , where the constant of proportionality depends only on d). If P is a singleton, $Z_P(\mathcal{S})$ is the cell of $\mathcal{A}(\mathcal{S})$ that contains P , so our analysis applies to single cells as well.

We obtain a new and simple proof of the following result; earlier proofs were given in [3] and [5]:

Theorem 5.1. *The maximum number of vertices in a single cell or in the zone of a fixed-degree algebraic surface, or of the boundary of any convex set, in an arrangement of n $(d - 1)$ -simplices in \mathbb{R}^d is $O(n^{d-1} \log n)$.*

Proof. Let $\overline{Z_P}(\mathcal{S})$ be the closure of the union of $Z_P(\mathcal{S})$. A k -face of $\mathcal{A}(\mathcal{S})$ is called *inner* if it lies in the intersection of the relative interiors of exactly $d - k$ simplices of \mathcal{S} and avoids simplex boundaries. A k -face is called *outer* if it is contained in the relative boundary of at least one simplex in \mathcal{S} . An inner k -face f is said to be *popular* if $f \subseteq \overline{Z_P}(\mathcal{S})$ but $f \not\subseteq \partial \overline{Z_P}(\mathcal{S})$. For instance, every cell of $Z_P(\mathcal{S})$ is popular, and a popular facet is one that touches cells of $Z_P(\mathcal{S})$ on both sides. This notation is borrowed from [5].

For each simplex $\sigma \in \mathcal{S}$, let h_σ be the hyperplane containing σ . For each inner vertex v of $\mathcal{A}(\mathcal{S})$, the triple (v, R, F) is called a *k-border* if the two following conditions are satisfied:

- R is a closed cone of \mathbb{R}^d whose interior is a connected component of $\mathbb{R}^d \setminus \bigcup_{i=1}^d h_{\sigma_i}$, where $\sigma_1, \dots, \sigma_d$ are the d simplices incident to v . Each inner vertex v induces 2^d cones R of this form.
- F is an intersection of $d - k$ hyperplanes from $\{h_{\sigma_1}, \dots, h_{\sigma_d}\}$ (F is a k -flat containing v). If $k = d$, then F must be \mathbb{R}^d .

Let (v, R, F) be any k -border and let f be the unique k -face $f \subset F$ of $\mathcal{A}(\mathcal{S})$, having v incident to $f \cap R$. We call f the face of (v, R, F) in $\mathcal{A}(\mathcal{S})$. If f is a popular face, we say that (v, R, F) is a *0-level popular k-border* of $Z_P(\mathcal{S})$.

For a simplex $\sigma \in \mathcal{S}$ we say that the quadruple (v', R', F', σ) is a *1-level popular k-border* if (v', R', F') is not a 0-level popular k -border of $Z_P(\mathcal{S})$, but it is a 0-level popular k -border of $Z_P(\mathcal{S} \setminus \{\sigma\})$. Let $C_i^{(k)}(P, \mathcal{S})$ be the number of i -level popular k -borders, for $i = 0, 1$ and $k = 0, 1, \dots, d$, and let $C_i^{(k)}(P, n)$ be the maximum of $C_i^{(k)}(P, \mathcal{S})$ over all sets \mathcal{S} of n $(d - 1)$ -simplices in \mathbb{R}^d in general position. (Note that $C_0^{(k)}(P, \mathcal{S})$ is equal to the quantity $\tau_0^{(k)}(p, \mathcal{S})$ in [5], when P is the single point p used in that paper, although the subscript 0 has a different meaning in each case.) The quantity $C_0^{(d)}(P, \mathcal{S})$ counts all the inner vertices on the boundary of every cell in $Z_P(\mathcal{S})$, where a vertex is multiply counted, once for every 0-level popular d -border that it generates within $Z_P(\mathcal{S})$. Since the number of all outer vertices in $\mathcal{A}(\mathcal{S})$ is $O(n^{d-1})$, it suffices to show that $C_0^{(d)}(P, \mathcal{S}) = O(n^{d-1} \log n)$.

Let $b_0 = (v_0, R_0, F)$ be a 0-level popular k -border of $Z_P(\mathcal{S})$, and let f_0 be the (popular) k -face of b_0 in $\mathcal{A}(\mathcal{S})$. Let e_1 be an edge of f_0 incident to v_0 and within R_0 , and let σ_1 be the (unique) simplex containing v_0 but not containing e_1 . Let e'_1 be the other edge in $\mathcal{A}(\mathcal{S})$ incident to v_0 but not contained in σ_1 (e_1 and e'_1 are adjacent edges on the same line), and let v_1 be the other endpoint of e'_1 . See Fig. 3 for an illustration.

If v_1 is an inner vertex, then let R_1 be the unique “cone” induced by v_1 and containing R_0 , and let f_1 be the face (whether popular or not) of (v_1, R_1, F) in $\mathcal{A}(\mathcal{S})$. Let g_1 be the $(k - 1)$ -face of $(v_0, R_0, F \cap h_{\sigma_1})$ in $\mathcal{A}(\mathcal{S})$. Note that if v_1 is an inner vertex, then e'_1 is incident to f_1 , $v_0 \in R_1$, g_1 is a $(k - 1)$ -face incident to f_0 and f_1 , and if we remove σ_1 from \mathcal{S} , then the faces f_0 and f_1 become part of a larger inner k -face f'_1 , which is clearly also popular. Note that it is possible that $f_0 = f_1$; this is reflected in the case




Fig. 3. Moving from an inner vertex v_0 to another vertex v_1 along an edge transversal an incident simplex.

analysis given below. We conclude that if v_1 is an inner vertex, then (v_1, R_1, F) is a 0-level popular k -border of $Z_P(\mathcal{S} \setminus \{\sigma\})$. In any case, one of the following three types of configurations must arise:

- (i) v_1 is an outer vertex.
- (ii) f_1 is a popular face (in $Z_P(\mathcal{S})$). In this case g_1 is also a popular face and $b' = (v_0, R_0, F \cap h_{\sigma_1})$ is a 0-level popular $(k - 1)$ -border of $Z_P(\mathcal{S})$.
- (iii) v_1 is an inner vertex and f_1 is not a popular face. Then $b_1 = (v_1, R_1, F, \sigma_1)$ is a 1-level popular k -border of $Z_P(\mathcal{S})$. We say that b_0 and b_1 are neighbors in $\mathcal{A}(\mathcal{S})$.

We repeat this analysis for each of the k edges of f_0 incident to v_0 and within R_0 .

For type (i) configurations, we charge the outer vertex v_1 by 2 units. The vertex v_1 is in the intersection of $d - 1$ simplices and on the boundary of one of these simplices. The only inner edge incident to v_1 is thus e'_1 , so v_1 can be charged in type (i) configurations only along e'_1 by 0-level popular k -borders of the form (v_0, R', F') , where e'_1 is contained by F' but not in R' . There are 2^{d-1} such cones R' and $\binom{d-1}{d-k}$ such k -flats F' . Thus each outer vertex can be charged at most $\binom{d-1}{d-k} 2^{d-1}$ times in type (i) configurations, and the total charge for this type of events is thus $O(n^{d-1})$, where the constant of proportionality depends on d .

For type (ii) configurations, we charge the popular $(k - 1)$ -border b' in the configuration by 2 units. A 0-level popular $(k - 1)$ -border (v_0, R_0, G) can be charged only by 0-level popular k -borders of the form (v_0, R_0, F') where $G \subset F'$. Since there are $d - k + 1$ such k -flats F' , it follows that (v_0, R_0, G) can be charged at most $d - k + 1$ times in type (ii) configurations, and the total charge for this type of events is at most $2(d - k + 1)C_0^{(k-1)}(P, \mathcal{S})$.

For type (iii) configurations, we charge the 1-level popular k -border b_1 ; we may also charge some outer vertices of $Z_P(\mathcal{S})$ (see case (b(i)) below). Let $m(b_1)$ denote the number of 0-level popular k -border neighbors of b_1 , which is the number of times b_1 will be charged (in type (iii) events). Since b_1 can be charged only along the k edges of f_1 incident to v_1 and within R_1 , we have $m(b_1) \leq k$. We give b_1 a total of 1 unit of charge, and it pays $1/m(b_1)$ units to each 0-level neighbor. Thus the charge it can pay to a 0-level neighbor will be at least $1/k$ units. This implies that any 0-level popular k -border receives at least 1 unit of charge (it will receive at least 2 units of charge if we reach a configuration of type (i) or (ii) along at least one of the k incident edges).

We want to show that, for $k \geq 3$, every 0-level popular k -border will receive more than 1 unit. As just noted, this is the case when at least one of the charging configurations




Fig. 4. Two 1-level neighbors on the same new simplex.

involving such a border is of type (i) or (ii), so we only consider popular k -borders for which all these charging configurations are of type (iii).

We denote each inner vertex of $\mathcal{A}(\mathcal{S})$ by the (unordered) tuple of the d simplices incident to the vertex. Let $b_0 = (v_0, R_0, F)$ be a 0-level popular k -border, with the inner vertex v_0 represented by the tuple $(1, 2, \dots, d)$. Let $b'_1 = (v'_1, R'_1, F, 1)$ be a 1-level popular k -border neighbor of b_0 , and suppose, without loss of generality, that v'_1 is represented by $(x, 2, \dots, d)$, for some new simplex $x \in \mathcal{S}$. Let $b'_2 = (v'_2, R'_2, F, 2)$ be another 1-level k -popular neighbor of b_0 , with v'_2 represented by $(1, y, 3, \dots, d)$, for a simplex $y \in \mathcal{S}$. We consider the following three cases:

- (a) $x = y$. In this case it is easily seen that b'_1 (and also b'_2) can have at most $k - 1$ 0-level neighbors. Indeed, the representation of (the vertex of) any such 0-level neighbor b is obtained by removing one of the simplices $x, 2, \dots, d$ incident to v'_1 and replacing it with the “hiding” simplex 1. Moreover, the vertex of b must be contained in R'_1 . However, one of the vertices obtained in this manner is v'_2 (when we replace simplex 2 by simplex 1), which is not contained in R'_1 , as easily follows by construction; see Fig. 4. It follows that b'_1 can pay at least $1/(k - 1)$ units of charge to b_0 , so b_0 gets a total of at least

$$\frac{1}{k - 1} + \frac{k - 1}{k} = 1 + \frac{1}{k(k - 1)} > 1$$

units of charge.

- (b) $x \neq y$. Suppose that b'_1 has a 0-level neighbor $b_1 = (v_1, R_1, F)$, and that b'_2 has a 0-level neighbor $b_2 = (v_2, R_2, F)$. We consider the following two subclasses:
 - (i) Suppose that v_1 and v_2 are both obtained by replacing the same simplex, say simplex 3, by the respective hiding simplices 1 and 2. That is, v_1 is represented by $(x, 1, 2, 4, \dots, d)$ and v_2 is represented by $(y, 1, 2, 4, \dots, d)$. It is easily seen by construction, that both v_1 and v_2 must belong to R_0 . Moreover, the three vertices all lie on the same line segment s , formed by the intersection of the $d - 1$ simplices $1, 2, 4, \dots, d$. Since v_0 is the apex of R_0 , it follows that v_0 cannot be the middle vertex along s . Suppose, with no loss of generality, that the middle vertex is v_1 , and consider the triangle τ with vertices v_0, v_2 , and v'_2 . See Fig. 5 for an illustration. Let Δ be the 2-flat containing the intersection of the simplices $1, 4, \dots, d$, and thus containing τ . By construction, the edges $v_0v'_2$ and $v_2v'_2$ of τ are not crossed by any simplex of \mathcal{S} . On the other hand, v_1 , on the edge v_0v_2 , is the intersection of s with the simplex x . We conclude that f'_2 , the 2-face of (v'_2, R'_2, Δ) in $\mathcal{A}(\mathcal{S})$, is not convex and thus has an outer vertex p' as one




Fig. 5. Charging an outer vertex in the triangle τ .

of its vertices. We charge p' by 4 units, to be divided between b_0 and b_2 . The vertex p' is contained in the intersection of Δ with a simplex, say σ , and is on the boundary of σ . It follows that p' can be charged in type (iii) configuration only as a vertex of a unique 2-face in Δ , namely the 2-face f'_2 . It also follows that p' can only be charged, via f'_2 , by the pair b_0 and b_2 , and this charging already implies that b_0 receives at least 2 units of charge in this case. Since the number of outer vertices is $O(n^{d-1})$, the total charge of this kind is also $O(n^{d-1})$.

- (ii) Suppose next that the preceding subcase does not occur, and that v_1 is not incident to simplex j , for some $j \in \{3, \dots, d\}$ (that is, v_1 is obtained by replacing simplex j by simplex 1 in the tuple defining v'_1). Then v_2 must be incident to simplex j , or else the preceding subcase would occur. This is easily seen to imply that $m(b'_1) + m(b'_2) \leq k + 2$, so at least one of $m(b'_1), m(b'_2)$ must be strictly smaller than k (for $k \geq 3$), which means that b_0 will receive, as in case (a), at least $1 + 1/k(k - 1)$ units of charge.

If we put $\rho_k = 1 + 1/k(k - 1)$, we can conclude that, for $k \geq 3$,

$$\rho_k C_0^{(k)}(P, \mathcal{S}) \leq C_1^{(k)}(P, \mathcal{S}) + O(n^{d-1} + C_0^{(k-1)}(P, \mathcal{S}))$$

and that, for $k = 2$,

$$C_0^{(2)}(P, \mathcal{S}) \leq C_1^{(2)}(P, \mathcal{S}) + O(n^{d-1} + C_0^{(1)}(P, \mathcal{S})).$$

Remark. Using arguments similar to those in the proof of Theorem 4.1, we can replace ρ_k by $\frac{4}{3}$, for $k \geq 3$.

We now apply our probabilistic analysis technique. Since 1-level borders are not quite the same as 0-level borders, we go over the application in more detail. That is, let $\mathcal{R} \subset \mathcal{S}$ be a random sample of $n - 1$ simplices. Then the expected number of 0-level popular k -borders of $Z_P(\mathcal{R})$ is still

$$\frac{n - d}{n} C_0^{(k)}(P, \mathcal{S}) + \frac{1}{n} C_1^{(k)}(P, \mathcal{S}).$$

This follows from the fact that the probability that a 1-level popular k -border of $Z_P(\mathcal{S})$ will yield a 0-level popular k -border of $Z_P(\mathcal{R})$ is $1/n$. More precisely, a triple (v, R, F) may have several simplices $\sigma \in \mathcal{S}$ such that the quadruple (v, R, F, σ) is a 1-level

popular k -border of $Z_P(\mathcal{S})$. When each of these simplices is removed, the triple (v, R, F) becomes a 0-level popular k -border of $Z_P(\mathcal{R})$. This is easily seen to imply the above equality. Thus, using our technique, we obtain, for $k \geq 3$, the recurrence

$$\frac{n-d+\rho_k}{n} C_0^{(k)}(P, n) \leq C_0^{(k)}(P, n-1) + \frac{1}{n} O(n^{d-1} + C_0^{(k-1)}(P, n)), \quad (9)$$

and, for $k = 2$,

$$\frac{n-d+1}{n} C_0^{(2)}(P, n) \leq C_0^{(2)}(P, n-1) + \frac{1}{n} O(n^{d-1} + C_0^{(1)}(P, n)). \quad (10)$$

As shown in [3] and [5], the number of cells in $Z_P(\mathcal{S})$ is $O(n^{d-1})$, where the constant of proportionality depends on d (and on the degree of P if P is an algebraic surface). By the Chopping Theorem of [5], the cells in $Z_P(\mathcal{S})$ can be divided into $O(n^{d-1})$ convex polyhedra with pairwise-disjoint interiors. If (v, R, F) is a 0-level popular 1-border, then v is either a locally lowest vertex or a locally highest vertex (relative to the x_d -direction) of a cell in $Z_P(\mathcal{S})$. By rotating the coordinate system slightly, we may assume that each of these convex polyhedra contains (at most) one locally lowest vertex and one locally highest vertex. Since any such vertex belongs to at most $d2^d$ popular 1-borders, we conclude that $C_0^{(1)}(P, n) = O(n^{d-1})$. Using the recurrence (10), we get, by Proposition 3.1, $C_0^{(2)}(P, n) = O(n^{d-1} \log n)$. Then, for $k \geq 3$, we proceed by induction on k , and use the recurrence (9). Applying again Proposition 3.1, we obtain $C_0^{(k)}(P, n) = O(n^{d-1} \log n)$, for $k = 3, \dots, d$. This completes the proof of the theorem. \square

6. Vertical Decompositions in Arrangements of Triangles in \mathbb{R}^3

In this section we obtain the main new result of the paper: improved bounds on the complexity of vertical decompositions in arrangements of triangles in 3-space, improving and extending previous bounds in [13].

Let \mathcal{T} be a set of n triangles in \mathbb{R}^3 . Let $\mathcal{A}(\mathcal{T})$ denote the arrangement of \mathcal{T} . Let P be a fixed point set in \mathbb{R}^3 . We denote by $Z_P(\mathcal{T})$ the zone of P in $\mathcal{A}(\mathcal{T})$, as defined in the previous section. Let e_1, e_2 be two edges of $\mathcal{A}(\mathcal{T})$ intersecting a common vertical line ℓ . The open segment $s \subset \ell$ between e_1 and e_2 is said to be a k -level visibility segment, and the triple (e_1, e_2, s) is said to be a k -level visibility configuration, if exactly k triangles of \mathcal{T} intersect s , and s intersects a cell in $Z_P(\mathcal{T})$. We distinguish between outer visibility configurations, where at least one of the edges e_1, e_2 is a portion of the boundary of some triangle, and inner configurations, where each of these edges is a portion of the intersection of the relative interior of two triangles. Denote by $C_k(P, \mathcal{T})$ (resp. $D_k(P, \mathcal{T})$) the number of inner (resp. outer) k -level visibility configurations. Denote by $C_k(P, n)$ (resp. $D_k(P, n)$) the maximum of $C_k(P, \mathcal{T})$ (resp. $D_k(P, \mathcal{T})$) over all sets \mathcal{T} of n triangles in general position.

We want to analyze the combinatorial complexity of the vertical decomposition of $Z_P(\mathcal{T})$. This is a decomposition of the cells of $Z_P(\mathcal{T})$ into subcells of constant complexity, obtained as follows. We take each edge e of each cell C of $Z_P(\mathcal{T})$, and extend from each point on e a vertical segment up and down into C until it meets the boundary of C

again. The resulting collection of vertical faces partitions each cell of $Z_P(\mathcal{T})$ into subcells, each of which has a unique top face and a unique bottom face, each being a portion of some triangle in \mathcal{T} . The complexity of each of these subcells need not be constant, so we further refine each subcell, slicing it by planes parallel to the y -axis, into prism-like subcells of constant complexity. We refer the reader to [13] for more details concerning vertical decompositions in arrangements of triangles. As shown there, the complexity of the vertical decomposition of $Z_P(\mathcal{T})$ is proportional to the sum of the number of 0-level visibility configurations and the number of vertices of $Z_P(\mathcal{T})$, so we concentrate on bounding the number $C_0(P, \mathcal{T}) + D_0(P, \mathcal{T})$ of 0-level visibility configurations.

We begin by estimating the number of outer visibility configurations. Let e be a boundary edge of a triangle in \mathcal{T} , and let h be the vertical strip spanned by e (it is the union of all vertical lines intersecting e). Let h^+, h^- denote, respectively, the portions of h that lie above and below e . We intersect h with all the other triangles in \mathcal{T} , thereby obtaining a two-dimensional arrangement of segments within h . Let $\mathcal{A}_e^+, \mathcal{A}_e^-$ denote, respectively, the portions of this arrangement within h^+ and within h^- . Each 0-level outer visibility configuration involving e can be identified with a vertex of either the lower envelope of the arrangement \mathcal{A}_e^+ or the upper envelope of the arrangement \mathcal{A}_e^- . Since the complexity of any such envelope is $O(n\alpha(n))$ [21], and since there are $3n$ boundary edges of the triangles of \mathcal{T} , the total number of 0-level outer visibility configurations is

$$D_0(P, n) = O(n^2\alpha(n)),$$

even if P is the entire 3-space. Similarly, each 1-level outer visibility configuration can be identified with a 1-level vertex in some arrangement \mathcal{A}_e^+ or \mathcal{A}_e^- . Using a simple variant of the Clarkson–Shor probabilistic technique [11], it is easily seen that we also have

$$D_1(P, n) = O(n^2\alpha(n)).$$

Consider next the (more involved) case of inner visibility configurations. We may assume, with no loss of generality, that the given triangles are in general position. Otherwise, if we perturb them slightly, so as to put them in general position, the number of inner visibilities can only increase, as easily checked.

Let $v = (e_1, e_2, s)$ be a 0-level inner visibility configuration, where e_1 (resp. e_2) is a portion of the intersection of two triangles t_1, t_2 (resp. t_3, t_4) of \mathcal{T} , and where e_1 is above e_2 . The visibility segment s lies fully within a single cell, which is necessarily a cell of $Z_P(\mathcal{T})$, so, in particular, e_1 and e_2 are edges of $Z_P(\mathcal{T})$. We slide a vertical segment s' , starting from s , so that the top endpoint of s' moves along e_1 and its bottom endpoint moves along t_3 , so that s' crosses t_4 . (Later we will repeat this sliding in all other directions too.) There is always a unique direction of movement of s' where these conditions are satisfied; if we move in the reverse direction, similar properties will hold, with t_3 and t_4 interchanged. See Fig. 6 for an illustration. We stop the sliding process as soon as we first encounter one of the following types of events:

- (i) s' reaches an endpoint of e_1 , which is either an outer or an inner vertex (in the notation of the preceding sections) of some cell of $Z_P(\mathcal{T})$.
- (ii) s' reaches a boundary edge e_3 of t_3 , at the 1-level outer visibility configuration (e_1, e_3, s') .




Fig. 6. Sliding a vertical segment from a 0-level visibility configuration.

- (iii) s' passes through a boundary of a triangle between t_3 and e_1 , at an appropriate 1-level or 0-level outer visibility configuration.
- (iv) s' reaches an intersection edge e' , between t_3 and a new triangle, at the corresponding 1-level inner visibility configuration (e_1, e', s') .

We apply this sliding process four times, twice by sliding along e_1 and along one of the triangles t_3, t_4 (as described above), and twice by sliding along e_2 and along one of the triangles t_1, t_2 , in a fully symmetric manner.

An important observation is that the sliding process can never reach a 0-level inner visibility configuration before any of these terminal events are encountered, so we can reach any terminal event at most once in each possible sliding direction.

In any of the above cases, we charge the appropriate terminal configuration or vertex by 1 unit. Thus each 0-level inner visibility configuration receives 4 units of charge.

Each vertex v of some cell C of $Z_P(\mathcal{T})$ can be charged at most six times in events of type (i) (along the at most six edges of C incident to v), and each (0-level or 1-level) outer visibility configuration can be charged at most twice in events of type (ii) or (iii). Hence, if we denote by $M(P, \mathcal{T})$ the complexity of $Z_P(\mathcal{T})$, we can bound the total amount of charge made in cases (i), (ii), and (iii) by $O(n^2\alpha(n) + M(P, \mathcal{T}))$.

A 1-level inner visibility configuration may be charged at type (iv) events up to four times. We give each such configuration 2 units of charge, so, if more than two charges are made to such a configuration, we need to pass the extra charges to other features. This is done as follows.

Let $v = (e_1, e_2, s)$ be a 1-level inner visibility configuration, with e_1 lying above e_2 . The visibility segment s stabs a single triangle, which we denote by t_x , at some face $f_x \subseteq t_x$ of $\mathcal{A}(\mathcal{T})$. Let t_1, t_2 be the two triangles whose intersection contains e_1 and let t_3, t_4 be the two triangles whose intersection contains e_2 . If v is charged more than twice, then it must be charged by at least one 0-level inner visibility configuration v_1 of the form (e_1, e, σ) and by at least one 0-level inner visibility configuration v_2 of the form (e', e_2, σ') . We say that the first charge is made by sliding toward v “from above,” and that the second charge is made by sliding toward v “from below.” Then e must be a portion of the intersection of t_x with either t_3 or t_4 , and e' must be a portion of the intersection of t_x with either t_1 or t_2 . With no loss of generality, assume that $e \subseteq t_x \cap t_3$ and that $e' \subseteq t_x \cap t_1$. It is important to notice that e and e' are edges of the face f_x , and that f_x is a (*popular*) face of $Z_P(\mathcal{T})$.

Construct a path $\pi = \pi(v)$ within f_x that connects the points $\sigma \cap t_x$ and $\sigma' \cap t_x$ by two straight links with $s \cap t_x$ as a common point. See Fig. 7 for an illustration. Since v_1 and v_2 both charge v , it follows by definition that no event of type (i), (ii), (iii), or (iv) can occur between v_1 and v or between v_2 and v (except at v itself). This implies that the




Fig. 7. Constructing a path from a 1-level visibility configuration; the dashed edges are the projections of e_1, e_2 onto t_x .

relative interior of π does not pass directly below or above any edge of $\mathcal{A}(\mathcal{T})$, except for e_1 and e_2 (that is, there does not exist a vertical segment whose relative interior meets no triangle of \mathcal{T} , which connects a point on π with a point on any other edge of $\mathcal{A}(\mathcal{T})$). Also, π does not intersect any triangle, except for lying on t_x .

Note that the path $\pi(v)$ is not necessarily unique. For example, if v is charged by four 0-level inner visibility configurations, then we obtain four pairs of these configurations, where one configuration in each pair charges v from above and one charges v from below. Each such pair induces a path like the path $\pi(v)$, and all four paths pass through the common point $s \cap f_x$. To avoid this technical issue, we construct for each v as above only one of the (up to four) possible paths $\pi(v)$, and pass to this path up to 2 units of charge, leaving v with only 2 units of charge.

We construct one such path for each 1-level inner visibility configuration v' whose vertical segment stabs f_x and which is charged at least once “from above” and at least once “from below,” as v was charged. This will produce a collection of paths within f_x , each connecting two points on its boundary. The above analysis is easily seen to imply that no pair of those paths can cross each other. We repeat this construction over all popular faces f_x of $Z_P(\mathcal{T})$.

Our next goal is to bound the number of paths $\pi(v)$. The system of these paths within a fixed popular face f_x can be regarded as a plane embedding of a planar graph $G(f_x)$, whose nodes are represented by the edges of f_x and whose edges are represented by the paths $\pi(v)$. By Euler’s formula, the number of such paths is bounded by three times the number of nodes (edges of f_x) plus the number of graph-faces of degree 2; each such face is bounded by two paths $\pi(v), \pi(v')$, connecting between the same pair of edges, e, e' of f_x , and by appropriate portions of the edges e and e' . We thus need to bound the number of these degree-2 faces. See Fig. 8 for an illustration.

Suppose there are $z + 1$ adjacent paths that connect between the same pair of edges e, e' of f_x and create z degree-2 faces between them, where $z \geq 2$ (see Fig. 8, where the case $z = 2$ is illustrated). Suppose that, as above, $e \subseteq t_x \cap t_3$ and $e' \subseteq t_x \cap t_1$. By construction, each of these paths must then start on e at a 0-level inner visibility configuration whose



Fig. 8. Constructing a planar graph from the paths $\pi(v)$; two faces of degree 2 are shaded.




Fig. 9. The vertical envelope induced by “parallel” paths starting from the edge e .

other edge lies on t_1 above e . In other words, consider the vertical strip h spanned by e and define in the upper portion of h the two-dimensional arrangement \mathcal{A}_e^+ , as above. Then the $z + 1$ paths above induce $z + 1$ inner vertices of the lower envelope of \mathcal{A}_e^+ that lie all on the segment $s_1 = t_1 \cap h$. However, as is easily seen, for any triple a, b, c of such vertices appearing in this order along s_1 , there exists an endpoint of some segment in \mathcal{A}_e^+ that appears on the lower envelope between a and c . See Fig. 9. It follows that there are at least $z' = \lfloor z/2 \rfloor$ such endpoints over the portion of e bounding our z degree-2 faces. Each of these endpoints induces a 0-level *outer* visibility configuration, and we can charge our z faces to these z' outer visibility configurations, concluding that the number of such faces is at most three times the number of 0-level outer visibility configurations involving edges of f_x . This analysis fails when $z = 1$, but in such a case we simply ignore the degree-2 face, replace the two adjacent paths bounding it by a single path, and note that none of the two faces adjacent to this path has degree 2. All this implies that the number of paths drawn within a face f_x is bounded by six times the number of edges of f_x plus three times the number of 0-level outer visibility configurations involving edges of f_x . Summing these bounds over all (popular) faces of $Z_P(\mathcal{T})$, we conclude that the overall amount of excess charges made to 1-level inner visibility configurations (over 2 units of charge per configuration) is $O(M(P, \mathcal{T}) + n^2\alpha(n))$. In other words, we have shown that

$$4C_0(P, \mathcal{T}) \leq 2C_1(P, \mathcal{T}) + O(M(P, \mathcal{T}) + n^2\alpha(n))$$

or

$$2C_0(P, \mathcal{T}) \leq C_1(P, \mathcal{T}) + O(M(P, \mathcal{T}) + n^2\alpha(n)). \tag{11}$$

We can now obtain our first main result:

Theorem 6.1. *Let P be an algebraic surface in \mathbb{R}^3 of dimension 0, 1, or 2 and of constant degree, or the boundary of an arbitrary convex set. Then the vertical decomposition of the zone of P in the arrangement $\mathcal{A}(\mathcal{T})$ of a set \mathcal{T} of n triangles in \mathbb{R}^3 consists of $O(n^2 \log^2 n)$ cells.*

Proof. Since the complexity of the zone of P is $O(n^2 \log n)$ (see [3] and [5] and Section 5), (11) becomes

$$2C_0(P, \mathcal{T}) \leq C_1(P, \mathcal{T}) + O(n^2 \log n).$$

We now plug this into our probabilistic analysis technique, and note that the abstract dimension of the problem is 4, because each inner visibility configuration is defined in

terms of four triangles. The recurrence (6) thus becomes

$$\frac{n-2}{n}C_0(P, n) \leq C_0(P, n-1) + \frac{1}{n}O(n^2 \log n),$$

whose solution is easily seen to be $C_0(P, n) = O(n^2 \log^2 n)$ (see Proposition 3.1). The asserted bound now follows from the fact, mentioned above, that the number of cells in the vertical decomposition is proportional to the sum of the complexity of $Z_P(T)$ and of $C_0(P, n) + D_0(P, n)$. \square

If P is the entire 3-space, we can modify the above analysis in the following fairly straightforward manner. Let $\mathcal{S} \subset \mathcal{T}$ be a random sample of m triangles. Let $K(m) = \mathbf{E}_{\mathcal{S}}(M(\mathbb{R}^3, \mathcal{S}))$, where $\mathbf{E}_{\mathcal{S}}$ denotes expectation with respect to the random sample \mathcal{S} . Every inner vertex of $\mathcal{A}(T)$ has probability $\binom{n-3}{m-3} / \binom{n}{m}$ to appear in $\mathcal{A}(\mathcal{S})$ (this will happen if and only if the three triangles incident to the vertex are chosen in \mathcal{S}), and the number of outer vertices in $\mathcal{A}(\mathcal{S})$ is $O(m^2)$. Thus

$$K(m) = O\left(\frac{m^3}{n^3}M(\mathbb{R}^3, T) + m^2\right).$$

Let $\overline{C}_k(m) = \mathbf{E}_{\mathcal{S}}(C_k(\mathbb{R}^3, \mathcal{S}))$, and let $\mathbf{E}_{\mathcal{R}}$ denote expectation with respect to a random sample $\mathcal{R} \subseteq \mathcal{S}$ of size $|\mathcal{R}| = m - 1$. Plugging (11) into our new technique, we get

$$\begin{aligned} \frac{m-2}{m}\overline{C}_0(m) &= \frac{m-2}{m}\binom{n}{m}^{-1} \sum_{\substack{\mathcal{S} \subset \mathcal{T} \\ |\mathcal{S}|=m}} C_0(\mathbb{R}^3, \mathcal{S}) \\ &\leq \binom{n}{m}^{-1} \sum_{\substack{\mathcal{S} \subset \mathcal{T} \\ |\mathcal{S}|=m}} \left(\mathbf{E}_{\mathcal{R}}(C_0(\mathbb{R}^3, \mathcal{R})) + \frac{1}{m}O(M(\mathbb{R}^3, \mathcal{S}) + m^2\alpha(m)) \right) \\ &= \overline{C}_0(m-1) + O\left(\frac{1}{m}K(m) + m\alpha(m)\right) \\ &= \overline{C}_0(m-1) + O\left(\frac{m^2}{n^3}M(\mathbb{R}^3, T) + m\alpha(m)\right). \end{aligned}$$

The solution of this recurrence is easily seen to be (again using Proposition 3.1)

$$\overline{C}_0(m) = O\left(\frac{m^3}{n^3}M(\mathbb{R}^3, T) + m^2\alpha(m) \log m\right).$$

Putting $m = n$, we thus conclude:

Theorem 6.2. *The vertical decomposition of the entire arrangement of a set of n triangles in \mathbb{R}^3 in general position consists of $O(K + n^2\alpha(n) \log n)$ cells, where K is the complexity of the arrangement.*

Remark. Theorems 6.1 and 6.2 improve previous bounds given in [13]. Those bounds were $O(n^{2+\varepsilon})$, for an $\varepsilon > 0$, for the cases covered in Theorem 6.1, and $O(K + n^{2+\varepsilon})$,

for any $\varepsilon > 0$, for the entire arrangement (where the constants of proportionality depend, rather badly, on ε). Using the algorithm given in [13] and Theorem 6.2, the vertical decomposition of the entire arrangement can be computed (deterministically) in $O(K \log n + n^2 \alpha(n) \log^2 n)$ time. Using the randomized algorithm given in [12], the vertical decomposition of the zone of P (as defined, for an appropriate set P , in Theorem 6.1) can be computed in $O(n^2 \log^3 n)$ expected time.

Finally, we consider cases where P designates more complex portions of $\mathcal{A}(\mathcal{T})$. By an appropriate adaptation of the preceding analysis, we obtain:

Theorem 6.3.

- (a) *If P is a connected curve that has L crossing points with the triangles of \mathcal{T} , then the complexity of the vertical decomposition of the zone of P is $O(N^{3/2} L^{1/2} + n^2 \log^2 n)$.*
- (b) *The complexity of the vertical decomposition of all the nonconvex cells of $\mathcal{A}(\mathcal{T})$ is $O(n^{5/2})$.*
- (c) *The complexity of the vertical decomposition of any m cells of $\mathcal{A}(\mathcal{T})$ is $O(m^{1/3} n^2 + n^2 \log^2 n)$.*

Remark. The bounds in the preceding theorem are significantly larger than the actual complexity of the relevant cells. For example, as shown in [1], the complexity of all nonconvex cells is $O(n^{7/3})$, and the complexity of any m cells is $O(m^{2/3} n + n^2 \log n)$. Nevertheless, all bounds states in the theorem appear to be new.

Proof. We start with the proof of (a). If we apply the analysis technique obtained above, we get the recurrence

$$\frac{n-2}{n} C_0(P, \mathcal{T}) \leq \mathbf{E}(C_0(P, \mathcal{R})) + O\left(\frac{1}{n} M(P, \mathcal{T}) + n\alpha(n)\right),$$

where P is the given curve, and where \mathcal{R} is a random sample of $n - 1$ triangles of \mathcal{T} . If we divide the recurrence by $(n - 1)(n - 2)$, and put $V(P, \mathcal{T}) = C_0(P, \mathcal{T})/|\mathcal{T}|(|\mathcal{T}| - 1)$, we obtain

$$V(P, \mathcal{T}) \leq \mathbf{E}(V(P, \mathcal{R})) + O\left(\frac{1}{n^3} M(P, \mathcal{T}) + \frac{\alpha(n)}{n}\right).$$

We unfold this recurrence, and replace $V(P, \mathcal{T})$ by the maximum value $V(P, n)$ of this quantity, over all collections of n triangles in \mathbb{R}^3 (with the same original curve P), to obtain

$$V(P, n) \leq \sum_{j=1}^n \frac{M(P, j)}{j^3} + \sum_{j=1}^n \frac{\alpha(j)}{j}, \tag{12}$$

where $M(P, j)$ is the expected complexity of the zone P in an arrangement of a random sample of j triangles of \mathcal{T} . We note that the expected number of crossings of P with the triangles in such a random sample is Lj/n , so the expected number of cells in the zone of P in the sample arrangement is also Lj/n . Using the bound in [1], we can easily obtain

$M(P, j) = O((Lj/n)^{2/3} j + j^2 \log j)$. (We use here the fact that $\mathbf{E}(x^{2/3}) \leq (\mathbf{E}(x))^{2/3}$.) We substitute this bound into (12), and rewrite the first sum in the bound as

$$\sum_{j=1}^{(L/n)^{1/2}} \frac{M(P, j)}{j^3} + \sum_{j=(L/n)^{1/2}+1}^n \frac{M(P, j)}{j^3}.$$

In the first subsum we simply use the trivial bound $M(P, j) = O(j^3)$, so this subsum is $O((L/n)^{1/2})$.

In the second subsum we use the bound on $M(P, j)$ derived above, to obtain

$$\begin{aligned} \sum_{j=(L/n)^{1/2}+1}^n O\left(\frac{(L/n)^{2/3}}{j^{4/3}} + \frac{\log j}{j}\right) &= O\left(\frac{(L/n)^{2/3}}{(L/n)^{1/6}} + \log^2 n\right) \\ &= O\left(\left(\frac{L}{n}\right)^{1/2} + \log^2 n\right). \end{aligned}$$

Since the second sum in (12) is only $O(\alpha(n) \log(n))$, it is subsumed in the bounds just obtained. This is easily seen to imply that $C_0(P, n) = o(L^{1/2} n^{3/2} + n^2 \log^2 n)$, thus proving (a).

Assertion (b) follows from (a) by taking P to be a spanning path of the union of the triangle boundaries; in this case $L = O(n^2)$.

We next prove (c). We apply a similar analysis, with P equal to a set of m “marking” points, one in each of the given cells. As shown in [1], we have $M(P, j) = O(m^{2/3} j + j^2 \log j)$. We obtain the same bound (12), and we write the first sum in that bound as

$$\sum_{j=1}^{m^{1/3}} \frac{M(P, j)}{j^3} + \sum_{j=m^{1/3}+1}^n \frac{M(P, j)}{j^3}.$$

The first subsum is $O(m^{1/3})$ (using, as above, the trivial estimate $M(P, j) = O(j^3)$). The second subsum is

$$\sum_{j=m^{1/3}+1}^n O\left(\frac{m^{2/3}}{j^2} + \frac{\log j}{j}\right) = O\left(\frac{m^{2/3}}{m^{1/3}} + \log^2 n\right).$$

Hence $C_0(P, T) = O((m^{1/3} + \log^2 n)n^2)$, as asserted. □

The bound in Theorem 6.3(b) can be used to obtain a lazy randomized incremental algorithm, of comparable efficiency, for constructing all nonconvex cells of $\mathcal{A}(T)$. This algorithm is based on the recent technique of [12], and is useful for translational motion planning in 3-space. In more detail, let B be an arbitrary polyhedron free to translate among a collection of polyhedral obstacles. We transform the problem as follows. We fix some reference point O on B , and represent each placement of B by the position of O . For each pair of a vertex (resp. edge, face) u of B and a face (resp. edge, vertex) v of some obstacle, we form the locus $t_{u,v}$ of all placements of O such that u touches v at the corresponding placements of B . Each $t_{u,v}$ is a planar polygon. We triangulate these polygons, and thus obtain a collection T of triangles in 3-space. Let Z_1, Z_2 be two

given placements of B (that is, of O). Clearly, any collision-free translational motion of B at which O moves from Z_1 to Z_2 cannot cross any triangle of \mathcal{T} , so such a motion is possible if and only if Z_1 and Z_2 lie in the same cell of $\mathcal{A}(\mathcal{T})$. To determine whether this is the case, we precompute all nonconvex cells of $\mathcal{A}(\mathcal{T})$, and process them for efficient point location. Then, given the pair of placements Z_1, Z_2 , we test whether Z_1 and Z_2 lie in nonconvex cells of $\mathcal{A}(\mathcal{T})$. If at least one of them lies in such a cell, then both must lie in the same nonconvex cell, so the point locations already determine whether a desired motion between Z_1 and Z_2 exists. Otherwise, Z_1 and Z_2 must lie in the same convex cell of $\mathcal{A}(\mathcal{T})$. It is easy to test whether this is the case: Take the straight segment Z_1Z_2 and test whether it intersects any triangle in \mathcal{T} . If not, we have found the desired motion (along Z_1Z_2); otherwise no such motion exists. We thus get an algorithm that requires about $O(n^{5/2})$ preprocessing time and storage, where n is the number of “contact” triangles, and can perform motion-planning queries of the above kind in $O(n)$ times per query.

Acknowledgment

I would like to thank Micha Sharir, my thesis adviser, for his help in preparing this paper. In particular, he has contributed some of the results in Theorem 6.3 and the discussion following it.

References

1. P. K. Agarwal and B. Aronov, Counting facets and incidences, *Discrete Comput. Geom.* 7 (1992), 359–369.
2. P. K. Agarwal, O. Schwarzkopf, and M. Sharir, The overlay of lower envelopes in three dimensions and its applications, *Discrete Comput. Geom.* 15 (1996), 1–13.
3. B. Aronov, M. Pellegrini, and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.* 9 (1993), 177–186.
4. B. Aronov and M. Sharir, Triangles in space, or: Building (and analyzing) castles in the air, *Combinatorica* 10(2) (1990), 137–173.
5. B. Aronov and M. Sharir, Castles in the air revisited, *Discrete Comput. Geom.* 12 (1994), 119–150.
6. B. Aronov and M. Sharir, On translational motion planning in three dimensions, *Proc. 10th ACM Symp. on Computational Geometry*, 1994, pp. 21–30. (To appear in *SIAM J. Comput.*)
7. B. Aronov and M. Sharir, The union of convex polyhedra in three dimensions, *Proc. 34th Ann. IEEE Symp. on Foundations of Computer Science*, 1993, pp. 518–527.
8. J. D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Proc. 11th ACM Symp. on Computational Geometry*, 1995, pp. 79–88.
9. L. P. Chew, K. Kedem, M. Sharir, B. Tagansky, and E. Welzl, Voronoi diagrams of lines in three dimensions under a polyhedral convex distance function, *Proc. 6th ACM–SIAM Symp. on Discrete Algorithms*, 1995, pp. 197–204. (To appear in *J. Algorithms*.)
10. K. Clarkson, New applications of random sampling in computational geometry, *Discrete Comput. Geom.* 2 (1987), 195–222.
11. K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.
12. M. de Berg, K. Dobrindt, and O. Schwarzkopf, On lazy randomized incremental construction, *Discrete Comput. Geom.* 14 (1995), 261–286.
13. M. de Berg, L. Guibas, and D. Halperin, Vertical decomposition for triangles in 3-space, *Discrete Comput. Geom.* 15 (1996), 35–62.

14. H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.
15. H. Edelsbrunner, The upper envelope of piecewise linear functions: Tight bounds on the number of faces, *Discrete Comput. Geom.* 4 (1989), 337–343.
16. H. Edelsbrunner, R. Seidel, and M. Sharir, On the zone theorem for hyperplane arrangements, *SIAM J. Comput.* 22 (1993), 418–429.
17. L. Guibas, D. Halperin, J. Matoušek, and M. Sharir, On vertical decomposition of arrangements of hyperplanes in four dimensions, *Discrete Comput. Geom.* 14 (1995), 113–122.
18. L. Guibas and M. Sharir, Combinatorics and algorithms of arrangements, in *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, New York, 1993, pp. 9–36.
19. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions with applications to visibility of terrains, *Discrete Comput. Geom.* 12 (1994), 313–326.
20. D. Halperin and M. Sharir, Arrangements and their applications in robotics: Recent developments, in *The Algorithmic Foundations of Robotics* (K. Goldberg, D. Halperin, J. C. Latombe, and R. Wilson, eds.), Peters, Boston, MA, 1995, pp. 495–511.
21. S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6 (1986), 151–177.
22. K. Mulmuley, *Computational Geometry: Introduction Through Randomized Algorithms*, Prentice-Hall, New York, 1993.
23. J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, *Discrete Comput. Geom.* 4 (1989), 291–309.
24. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.
25. M. Sharir and P. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge, 1995.
26. A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport–Schinzel sequences by segments, *Discrete Comput. Geom.* 3 (1988), 15–47.

Received April 26, 1995, and in revised form November 5, 1995, and March 17, 1996.