

Almost Tight Upper Bounds for Lower Envelopes in Higher Dimensions*

M. Sharir

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

and

Courant Institute of Mathematical Sciences, New York University,
251 Mercer Street, New York, NY 10012, USA

Abstract. We consider the problem of bounding the combinatorial complexity of the lower envelope of n surfaces or surface patches in d -space ($d \geq 3$), all algebraic of constant degree, and bounded by algebraic surfaces of constant degree. We show that the complexity of the lower envelope of n such surface patches is $O(n^{d-1+\epsilon})$, for any $\epsilon > 0$; the constant of proportionality depends on ϵ , on d , on s , the maximum number of intersections among any d -tuple of the given surfaces, and on the shape and degree of the surface patches and of their boundaries. This is the first nontrivial general upper bound for this problem, and it almost establishes a long-standing conjecture that the complexity of the envelope is $O(n^{d-2}\lambda_q(n))$ for some constant q depending on the shape and degree of the surfaces (where $\lambda_q(n)$ is the maximum length of (n, q) Davenport–Schinzel sequences). We also present a randomized algorithm for computing the envelope in three dimensions, with expected running time $O(n^{2+\epsilon})$, and give several applications of the new bounds.

1. Introduction

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of n surfaces or surface patches in d -space, for $d \geq 3$. This paper continues the study initiated in the companion paper [18], and addresses the problem of bounding the *combinatorial complexity* of the

*Work on this paper has been supported by NSF Grant CCR-91-22103, and by grants from the U.S.–Israeli Binational Science Foundation, the G.I.F., the German–Israeli Foundation for Scientific Research and Development, and the Fund for Basic Research administered by the Israeli Academy of Sciences.

lower envelope (pointwise minimum) of the surface patches in Σ (see below for a more precise definition of this term). Our analysis applies to collections of general algebraic surfaces. For example, in three dimensions our results hold for collections Σ for which all the surface patches in Σ are algebraic of some constant degree, and the boundary of each patch consists of a constant number of algebraic arcs, all of constant degree as well. We also assume, with no loss of generality, that each σ_i is the graph of a partially defined (algebraic) function $z = f_i(x, y)$; otherwise, we can cut each patch into subpatches along the (constant number of) curves of points of z -vertical tangency. We also assume that the given patches are in *general position*, meaning that the coefficients of the polynomials defining the surfaces and their boundaries are algebraically independent over the rationals. This assumption excludes degenerate configurations; for example, in three dimensions, it excludes configurations where four surfaces meet at a point, the boundary of one surface meets a curve of intersection of two other surfaces, two boundary curves of distinct surfaces meet at a point, etc. This involves no real loss of generality, because, as is argued below, the maximum complexity of the envelope of n surface patches with a fixed maximum algebraic degree and a fixed maximum degree of their boundaries, is at most proportional to the same maximum taken over collections of surfaces that are in general position. Note that the extra condition assumed in [18] (for the three-dimensional case)—that the relative interiors of any triple of surfaces intersect in at most two points—is no longer assumed.

In three dimensions the lower envelope E_Σ of Σ , its xy -projection \mathcal{M}_Σ , and their combinatorial complexity $\psi(\Sigma)$ are defined exactly as in [18]. These notions are also appropriately generalized to higher dimensions—see Section 2.

Under the assumptions made above, it is easy to show that, in d -space, $\psi(\Sigma) = O(n^d)$ (with a constant of proportionality that depends on d and on the algebraic degree of the patches and of their boundaries). However, it has been conjectured over the past 8 years that the maximum possible complexity of such an envelope is at most only slightly larger than $O(n^{d-1})$. This conjecture has previously been proven only for families of a few special types of surfaces or surface patches, such as hyperplanes, balls, simplices, and, in three dimensions, also for a few other types (see [26] and [28]).

As noted in [18], the conjecture was motivated by the fact that in two dimensions, in the case of the lower envelope of n partially defined univariate functions, the complexity of their envelope is at most $\lambda_{s+2}(n)$, which is the maximum length of *Davenport–Schinzel sequences* of order $s + 2$ composed of n symbols (see [4] and [21] for more details); here s is the maximum number of intersection points between any pair of arcs. Unfortunately, an appropriate extension of these results to three and higher dimensions (namely, the above-mentioned conjecture, whose precise form states that the complexity of the envelope is at most $O(n^{d-2}\lambda_s(n))$ for some constant s depending on the given surfaces; see [26]) appears to be extremely difficult, and the problem in general has been wide open ever since; in fact, even in three dimensions no general bounds better than $O(n^3)$ were known so far.

In this paper we extend the analysis technique of [18] to obtain the first general

nontrivial bounds for the complexity $\psi(\Sigma)$, and, moreover, get very close to establishing the above conjecture. More precisely, we show that $\psi(\Sigma) = O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$, with the constant of proportionality depending on ε , on d , on the maximum number s of intersection points between any d -tuple of surfaces in Σ , and also on the degree and shape of these surfaces and of their boundaries¹ (the parameter s is singled out here because the dependence of the constant on s is more significant than its dependence on the other parameters—see below). The proof is not too difficult; it uses induction on d , and extends the random-sampling technique of [18] in several ways, so as to make it apply to the more general situations studied here. Our bounds still leave a small gap from the conjectured envelope complexity, but otherwise constitute a rather satisfactory solution to the lower envelope problem in arbitrary dimensions.

Improved bounds on the complexity of lower envelopes have numerous applications in computational and combinatorial geometry. We present below some specific applications, where the new bounds yield improved algorithmic or combinatorial solutions to several geometric problems. Additional applications and properties of our results have recently been obtained, after the original preparation of this paper, and we mention some of them below. A comprehensive summary of geometric applications of lower envelopes is given in [30].

The paper is organized as follows. In Section 2 we prove the upper bound for the complexity of the envelope of surface patches in higher dimensions. Section 3 presents an efficient randomized algorithm for constructing the envelope in three dimensions, whose expected running time is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Section 4 presents some applications of our new bounds, and Section 5 concludes the paper with some comments and open problems.

2. Lower Envelopes of Surface Patches

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of $n(d-1)$ -dimensional surfaces or surface patches in d -space. We assume that these surfaces satisfy the following conditions, which are an appropriate generalization of the conditions assumed in [18] (we denote the coordinates by x_1, \dots, x_d , where x_d is the direction relative to which the lower envelope is defined):

- (i) Each σ_i is monotone in the $x_1x_2 \cdots x_{d-1}$ -direction (that is, every line parallel to the x_d -axis intersects σ_i in at most one point). Moreover, each σ_i is a portion of a $((d-1)$ -dimensional) algebraic surface of constant degree b .

¹ The term “shape” refers to additional properties that the given surfaces might possess, which restrict the patterns of intersections and other interactions between the surfaces, more than what follows only from the fact that the surfaces are algebraic of constant degree. For example, the shape of the surfaces might imply that the maximum number s of intersections between any d of the surfaces is small.

- (ii) The projection in the x_d -direction of σ_i onto the hyperplane $x_d = 0$ is a semialgebraic set defined in terms of a constant number of $(d - 1)$ -variate polynomials of constant degree (say, b too).
- (iii) The relative interiors of any d of the given surfaces intersect in at most s points, for some constant parameter s (by Bezout's theorem [22] and by property (iv) below, we always have $s \leq b^d$).
- (iv) The surface patches in Σ are in *general position*, in the sense defined in the Introduction.

The first part of condition (i) is not essential, since we can always cut a given algebraic surface of constant degree into a constant number of monotone pieces, so that condition (ii) continues to hold for these pieces. Condition (iv) is also not essential, as is argued below, and condition (iii) is made only for convenience, allowing us to carry out our analysis in terms of s , which is usually much smaller than b^d ; this results in significantly smaller constants of proportionality.

The lower envelope E_Σ of Σ is the graph of the (partial) function $x_d = E_\Sigma(x_1, \dots, x_{d-1})$ that maps each point (x_1, \dots, x_{d-1}) to the smallest x_d -coordinate among those of the points of intersection between the x_d -parallel line through (x_1, \dots, x_{d-1}) and the surfaces in Σ (if no such surface exists, the envelope is undefined at (x_1, \dots, x_{d-1})). If that lowest point lies on the boundary of one or more surfaces, we take the maximal closed segment contained in the x_d -parallel line through (x_1, \dots, x_{d-1}) whose bottom endpoint lies on the envelope and which does not cross the relative interior of any surface in Σ , and say that the envelope is *attained* over (x_1, \dots, x_{d-1}) by all surfaces that touch that segment; if the envelope point does not lie on any surface boundary, we say that the envelope is attained by (the relative interior of) all surfaces incident to that point. If we project E_Σ onto the hyperplane $H: x_d = 0$, we obtain a decomposition $\mathcal{M} = \mathcal{M}_\Sigma$ of H into connected relatively open semialgebraic sets, which we call *cells*, such that each cell c of \mathcal{M} is a maximal connected portion of H over which E_Σ is attained by a fixed combination of the relative interiors of some surfaces and/or the boundaries of other surfaces (by the general position assumption, the number of such surfaces is at most d). The *combinatorial complexity* of E_Σ , denoted by $\psi(\Sigma)$, is defined to be the number of cells (of all dimensions) of \mathcal{M} .

2.1. The Issue of General Position

As claimed in the Introduction, assuming that the surfaces in Σ are in general position involves no real loss of generality. This is because, roughly speaking, the complexity of the envelope of n surface patches with a fixed maximum algebraic degree and a fixed maximum degree of their boundaries, attains its maximum when the surfaces are in general position. Here is a more precise way of stating and proving this claim:

Suppose Σ is not in general position, and consider a slight perturbation of the

surfaces in Σ , obtained by replacing each polynomial $P_i(x_1, \dots, x_d) = 0$ defining these surfaces and their boundaries by the polynomial

$$P_i((1 - \epsilon_{i1}x_1 + \epsilon_{i1}), \dots, (1 - \epsilon_{id}x_d + \epsilon_{id})) = 0,$$

where the ϵ_{ij} 's are arbitrarily small algebraically independent transcendentals. Let Σ' denote the resulting collection of perturbed surfaces. If v is an *inner vertex* of E_Σ , namely, a vertex formed by the intersection of the relative interiors of d or more surfaces of Σ , and, furthermore, if these surfaces intersect *transversally* at v , then, as is easily checked, v will be replaced in $E_{\Sigma'}$ by one or more vertices (provided the ϵ_{ij} 's are sufficiently small, which we assume anyway). Thus the number of such vertices can only grow when we pass from Σ to Σ' .

There are several other kinds of vertices for which this argument does not apply:

- (a) *Boundary vertices*, namely, vertices formed by the intersection of a boundary of one surface with other surfaces. Clearly, any such vertex is determined by at most $d - 2$ additional surfaces, so the number of these vertices is clearly only $O(n^{d-1})$.
- (b) *Boundary-visible vertices*, namely, vertices v for which the bottom endpoint of the vertical segment defining v as above, lies on the boundary of some surface of Σ . Nevertheless, we have:

Lemma 2.1. *The number of boundary-visible vertices of E_Σ is $O(n^{d-1+\epsilon})$, for any $\epsilon > 0$, where the constant of proportionality depends on ϵ , d , and b . In three dimensions the bound improves to $O(n\lambda_q(n))$, for some constant q depending on the degree and shape of the given surfaces and of their boundaries.*

Proof. For the proof of the lemma, we need to assume inductively that our main result, Theorem 2.2, holds for all $d' < d$; when $d = 3$ this assumption is not needed—see below. Let $\sigma \in \Sigma$, let c denote its boundary, and let V be the vertical surface formed by the union of all vertical rays whose bottom endpoints lie in c . For each surface $\sigma_i \neq \sigma$ in Σ , let $\zeta_i = \sigma_i \cap V$. The properties that the surfaces satisfy imply that each ζ_i is an algebraic surface patch, of maximum dimension $d - 2$, whose boundary satisfies condition (ii) above. (Since we are not assuming general position, the dimension of ζ_i can be smaller than $d - 2$; however, the foregoing analysis applies to these degenerate configurations as well.) It easily follows that the surfaces ζ_i satisfy conditions (i)–(iii) above (note that the maximum degree of ζ_i can be much larger than b , but is still a constant). Note that all boundary-visible vertices w that lie over c also appear as vertices of the lower envelope, within V , of the surface patches ζ_i . By the results of [7] and [13], c can be decomposed into a constant number of subcells, each diffeomorphic to an open ball of some dimension. Thus, for each subcell c' , we can parametrize c' using $t \leq d - 2$ real parameters, so as to embed c' as an open ball in t -dimensional euclidean space. The portion of V lying above c' is then embedded in $(t + 1)$ -dimensional space, so that conditions (i)–(iii) continue to hold for the embedded surfaces ζ_i . This implies, using the induction hypothesis, that the complexity of the lower envelope

of the ζ_i 's over c' is $O(n^{1+\varepsilon}) = O(n^{d-2+\varepsilon})$, for any $\varepsilon > 0$. Summing up these bounds, over all subcells c' of c and over the boundaries c of all the surfaces of Σ , we conclude that the overall number of boundary-visible vertices of E_Σ is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$. In three dimensions the ζ_i 's are $O(n)$ algebraic arcs, each pair of which intersects in at most some constant number, q' , of points (or connected arcs). It follows that the number of vertices of their envelope is at most $O(\lambda_{q'+2}(n))$. Summing over all surfaces σ , we conclude that the number of boundary-visible vertices is $O(n\lambda_q(n))$, where the constant q is the maximum, over all surface boundaries, of the numbers $q' + 2$, where q' is as defined above. The arguments used in the three-dimensional case resemble the analysis techniques of [7] and [11]. \square

- (c) *Singular inner vertices*, namely, inner vertices that are singular² on at least one of their incident surfaces. Since the locus of singular points on each surface of Σ is an algebraic surface of dimension $\leq d - 2$ (and of constant degree), it can be argued, as in case (a) above, that the total number of such vertices is $O(n^{d-1})$.
- (d) *Nontransversal inner vertices*, namely, inner vertices v incident to d (or more) surfaces, such that v is a nonsingular point on each of these surfaces, but the d (or more) tangent hyperplanes to the surfaces at v are not linearly independent. (We can ignore degenerate cases where v is a point of intersection between fewer than d surfaces, because the number of such vertices is again only $O(n^{d-1})$.) The difficulty in this case is that a small perturbation of these surfaces in the “wrong” direction may cause v to disappear. (Think, e.g., of the case where, in three dimensions, v is a point of tangency between one surface, σ_1 , and the curve of intersection of two other surfaces, σ_2, σ_3 ; moving σ_1 away from this curve may cause v to disappear.) However, it is easily checked that if we choose the ϵ_{ij} 's randomly and independently from a sufficiently small neighborhood of 0, there is a constant probability (depending on d) that v will show up on the perturbed envelope as one or more vertices. This implies that an appropriate random choice of the ϵ_{ij} 's will cause at least a fixed fraction of these vertices to show up on the perturbed envelope.

To sum up, an appropriate random perturbation of the surfaces will result in a collection of surfaces in general position, and the (expected) number N' of vertices on its envelope satisfies $N' \geq cN - K$, where c is some absolute positive fraction (depending on d), N is the number of vertices on the original envelope, and $K = O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$ (or, in three dimensions, $K = O(n\lambda_q(n))$, for some constant q). Since the bounds that we will establish will be asymptotically larger than K , this argument implies that it suffices to establish our bounds for collections of surfaces, in general position, since these bounds will also apply, up to some constant factor, in degenerate configurations. (The complexity of the envelope also counts features other than vertices. However, in general position, a straightforward

² A point v is *singular* on an algebraic surface, defined by $P = 0$ for some polynomial P , if all the partial derivatives $\partial P / \partial x_i$, for $i = 1, \dots, d$, vanish at v ; see [5] and [22].

charging scheme shows that the total number of other features of E_Σ is proportional to the number of vertices, plus an “overhead” term, which is at most $O(n^{d-1})$. This argument can also be extended to degenerate configurations.)

2.2. *The Complexity of the Envelope*

Theorem 2.2. *The combinatorial complexity of the lower envelope of a collection of n $(d - 1)$ -dimensional surface patches in d -space that satisfy conditions (i)–(iii) is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε , d , s , and on the maximum degree b and the shape of the given surfaces and of their boundaries.*

Proof. The proof is based on induction on d . The base case $d = 2$ is immediate from standard Davenport–Schinzel theory [21]. Let $d \geq 3$ and suppose that the theorem is true for all $d' < d$. Let Σ be a collection of n $(d - 1)$ -dimensional surface patches in d -space, which satisfy the above conditions and are in general position. Let $\varphi(\Sigma)$ and denote the number of inner vertices of E_Σ . Thus we temporarily ignore boundary vertices and boundary-visible vertices, as defined in Section 2.1. This will not affect our bounds, because, as noted above, the number of boundary and boundary-visible vertices on E_Σ is only $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$. We also denote by $\varphi(n)$ the maximum of $\varphi(\Sigma)$, taken over all collections Σ of n surface patches that satisfy conditions (i)–(iv), with the same b and s . Finally, we denote by $\varphi^*(\Sigma)$ the total number of vertices of E_Σ of all types (including boundary vertices and boundary-visible vertices, as enumerated in Section 2.1), and by $\varphi^*(n)$ the maximum of $\varphi^*(\Sigma)$, taken over all collections Σ as above. The discussion in Section 2.1 easily implies that $\varphi^*(n) = \varphi(n) + O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.

As in [18], we extend each surface $\sigma \in \Sigma$ to a surface σ^+ by erecting an upward-directed vertical ray from each point on the boundary of σ . Let Σ^+ denote the collection of these extended surfaces, and let $\mathcal{A}(\Sigma^+)$ denote their arrangement.

Let p be an inner vertex of E_Σ , formed by the intersection of d surfaces $\sigma_1, \dots, \sigma_d \in \Sigma$. We may assume that p is a nonsingular point on each of these surfaces, and that the surfaces meet transversally at p (by the analysis in the preceding subsection, the number of all other inner vertices is only $O(n^{d-1})$). For $i = 1, \dots, d$, we denote by γ_i the curve of intersection between all the surfaces $\sigma_1, \dots, \sigma_d$, except for σ_i . Let $\beta_i(p)$ denote the maximal open connected x_i -monotone subarc of γ_i which has p as an endpoint and is disjoint from E_Σ . Note that, since each γ_i is algebraic of constant degree, γ_i can be partitioned into a constant number of connected x_i -monotone subarcs, so that each $\beta_i(p)$ is fully contained in one of these subarcs; the total number of such subarcs, over all intersection curves of the surfaces, is thus $O(n^{d-1})$.

Since p is a point of transversal intersection between these d surfaces, and since none of these surfaces is singular at p , these surfaces are smooth at p . Hence E_Σ can be approximated, at a neighborhood of p , by the lower envelope of the d tangent hyperplanes to these surfaces at p , which is easily seen to imply that the

d tangent directions of $\beta_1(p), \dots, \beta_d(p)$ at p , all pointing from p along their respective arcs, are such that the positive span of their orthogonal projections onto H is the entire hyperplane H . Thus at least one of these arcs, say $\beta_1(p)$, emanates from p in the positive x_1 -direction, and at least one arc emanates in the negative x_1 -direction.

We define the *index* $j = j(p)$ of p to be the number of points of intersection of $\sigma_1, \dots, \sigma_d$ that lie in the half-space $x_1 > x_1(p)$; by property (iii), j can range from 0 to $s - 1$. We define $\varphi^{(j)}(\Sigma)$, for $j = 0, \dots, s - 1$, to be the number of inner vertices of E_Σ whose index is at most j . We derive a recurrence relationship for $\varphi(n)$, by bounding each of the functions $\varphi^{(j)}$ in terms of $\varphi^{(j-1)}$ (with special handling of $\varphi^{(0)}$); the solution to this recurrence yields the asserted bounds.

Let p be an inner vertex as above, whose index is at most j . Let u denote the other endpoint of $\beta_1(p)$. Several cases can arise:

(a) u is an endpoint of a maximal connected x_1 -montone piece of the whole intersection curve γ_1 . In this case we charge p to u . Since each such u can be charged in this manner at most a constant number of times, and since there are only $O(n^{d-1})$ such points, over the entire collection of intersection curves, the number of vertices p of this kind is $O(n^{d-1})$.

(b) u is a point on E_Σ , and the excluded surface σ_1 does not intersect $\beta_1(p)$ and is not incident to u . Thus, if we start tracing $\beta_1(p)$ from p , we move initially along points hidden from E_Σ by σ_1 , which thus lies below the curve. It follows that $\beta_1(p)$ must contain a point w that lies directly above the boundary of σ_1 , and we charge p to the point w .

Let w' be the point of $\partial\sigma_1$ that lies directly below w , and let t denote the number of surfaces that either cross the portion $\delta_1 = pw$ of $\beta_1(p)$ or have a point on their boundary which lies directly below δ_1 . We fix some threshold parameter $k = k_j$, to be determined below, and consider the following two subclasses:

(i) $t > k$. Let Σ^+ denote the collection of the vertically extended surfaces, as defined above, and let $\mathcal{A}(\Sigma^+)$ denote their arrangement. By assumption, δ_1 contains at least t vertices of $\mathcal{A}(\Sigma^+)$, and we charge p to the block of the first k of these vertices in their order from p to w along δ_1 . Clearly, each vertex of $\mathcal{A}(\Sigma^+)$ can be charged in this manner only a constant number of times: By the general position assumption, any such vertex w lies on only d curves of $(d - 1)$ -wise intersection of the surfaces of Σ^+ . Along any of these curves, γ , the vertex w can be charged at most once, since the charging is made by a vertex p that lies on $\gamma \cap E_\Sigma$, and the portion of γ between p and w is disjoint from E_Σ .

Define the *level*, with respect to $\mathcal{A}(\Sigma^+)$, of a point ζ in d -space to be the number of surfaces of Σ^+ that lie strictly below ζ (which is the same as the number of original surfaces in Σ that lie below ζ); note that the vertices of E_Σ have level 0. Each of the charged vertices along δ_1 is easily seen to be at level at most k : the initial point p is at level 0, and the level of each vertex along δ_1 differs from the level of the preceding vertex by only ± 1 .

We next obtain an upper bound for the number of vertices of $\mathcal{A}(\Sigma^+)$ that lie

at level $\leq k$. For this we apply the random-sampling analysis technique of [12] and [29], as also used in [18]. That is, we choose a random sample \mathcal{R} of $r = n/k$ surfaces of Σ^+ , and construct their arrangement $\mathcal{A}(\mathcal{R})$. Let v be a vertex of $\mathcal{A}(\Sigma^+)$ at level $t \leq k$. The probability that v shows up as a vertex of $E_{\mathcal{R}}$ is

$$\binom{n-t-d}{r-d} / \binom{n}{r}$$

out of the $\binom{n}{r}$ possible samples \mathcal{R} , the samples for which v is a vertex of their lower envelope are precisely those that include the d surfaces meeting at v and do not include any of the t surfaces lying below v . Hence we have

$$\sum_{t=0}^k \frac{\binom{n-t-d}{r-d}}{\binom{n}{r}} F_t \leq \varphi^*(\mathcal{R}) \leq \varphi^*(r),$$

where F_t is the number of vertices v of $\mathcal{A}(\Sigma^+)$ at level t . Using standard calculations [12], [18], [29], we can conclude that, when substituting $r = n/k$,

$$\sum_{t=0}^k F_t = O\left(k^d \varphi^*\left(\frac{n}{k}\right)\right); \tag{1}$$

in other words, the number of vertices of $\mathcal{A}(\Sigma^+)$ at level $\leq k$ is $O(k^d \varphi^*(n/k))$, which in turn implies that the number of inner vertices p of E_{Σ} in this subcase is $O(k^{d-1} \varphi^*(n/k))$.

(ii) $t \leq k$. Let c denote the boundary of σ_1 , and let V_c denote the vertical surface formed by the union of all x_d -parallel rays whose bottom endpoints lie in c . For each surface $\sigma_i \neq \sigma_1$ in Σ , let $\zeta_i = \sigma_i \cap V_c$. As argued in the proof of Lemma 2.1, an appropriate decomposition of c into “simple” subcells and an appropriate parametrization of each subcell of c allow us to regard the surfaces ζ_i as surfaces in a $(q + 1)$ -dimensional euclidean space, for some $q \leq d - 2$, so that they satisfy conditions (i)–(iv), for appropriate constant degree and constant maximum number of $(q + 1)$ -wise intersection points. Moreover, w is a vertex of some subarrangement $\mathcal{A}_{c'}$ of the surfaces ζ_i , over some subcell c' of c , whose level is $\leq k$ (where the level of a point in this cross-sectional arrangement is defined as above).

We can clearly apply the bound (1), which summarizes the preceding analysis, to this case, to conclude that the number of vertices of $\mathcal{A}_{c'}$ at level $\leq k$ is $O(k^{q+1} \Phi_{q+1}(n/k))$, where $\Phi_{q+1}(m)$ denotes the maximum complexity of the lower envelope of m surfaces ζ_i , in $(q + 1)$ -dimensional space, which satisfy the above conditions. The induction hypothesis implies that $\Phi_{q+1}(m) = O(m^{q+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε, q , the maximum degree

of the surfaces ζ_i and the maximum number of $(q + 1)$ -wise intersection points. Hence, the number of vertices of $\mathcal{A}_{c'}$ at level $\leq k$ is

$$O\left(k^{q+1} \binom{n}{k}^{q+\varepsilon}\right) = O(k^{1-\varepsilon} n^{q+\varepsilon}) = O(k^{1-\varepsilon} n^{d-2+\varepsilon}).$$

Summing this bound over all subcells c' of c , and over the boundaries c of all surfaces of Σ , we conclude that the overall number of vertices p that are charged in this subcase is $O(k^{1-\varepsilon} n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.

(c) u is a point on E_Σ , and the excluded surface σ_1 does intersect the open arc $\beta_1(p)$ or is incident to its endpoint u . Let w be the nearest point to p along $\beta_1(p)$ where σ_1 meets this curve, or $w = u$ if $\beta_1(p)$ does not meet σ_1 . Let δ_1 denote the portion pw of $\beta_1(p)$, and let t denote the number of extended surfaces of Σ^+ that cross δ_1 . We consider the following two subcases:

(i) $t > k$. As in case (b(i)) above, we conclude that δ_1 contains at least k vertices of $\mathcal{A}(\Sigma^+)$, and we charge p to the first k of these vertices, in their order along δ_1 from p to w . As argued above, all these vertices are at level at most k in $\mathcal{A}(\Sigma^+)$, and each of them can be charged in this way only a constant number of times. Repeating the random-sampling argument given above, we conclude that the number of inner vertices p that are charged in this manner is $O(k^{d-1} \varphi^*(n/k))$.

(ii) $t \leq k$. Here we observe, arguing as above, that w lies at level at most t in $\mathcal{A}(\Sigma^+)$. Hence, if we remove the $t' \leq t$ surfaces that hide w from E_Σ , w becomes an inner vertex of the envelope of the remaining surfaces. Moreover, the index of w , by definition, is at most $j - 1$. To exploit this observation, we draw, as above, a random sample \mathcal{R} of $r = n/k$ surfaces from Σ . By definition, the expected number of inner vertices of $E_{\mathcal{R}}$ with index at most $j - 1$ is $\varphi^{(j-1)}(\mathcal{R}) \leq \varphi^{(j-1)}(r)$. Now, arguing as above, the probability that our w will show up as such a vertex of $E_{\mathcal{R}}$ is $\binom{n-t'-d}{r-d} / \binom{n}{r}$, so we obtain

$$\sum_{t=0}^k \frac{\binom{n-t-d}{r-d}}{\binom{n}{r}} G_t \leq \varphi^{(j-1)}(r),$$

where G_t is the number of inner vertices w of $\mathcal{A}(\Sigma)$ that arise in the present analysis, with exactly t surfaces passing below them. As above, this can be rewritten as

$$\sum_{t=0}^k G_t = O\left(k^d \varphi^{(j-1)}\left(\frac{n}{k}\right)\right);$$

in other words, if we charge each of the inner vertices p in this subcase to the corresponding point w , and observe that each such point can be charged only a constant number of times, we conclude that the number of vertices p in this subcase is $O(k^d \varphi^{(j-1)}(n/k))$.

Thus, summing up these bounds, we obtain the following recurrence: for any $\varepsilon' > 0$ there is a constant c (depending on ε' , d , and b) such that

$$\varphi^{(j)}(\Sigma) \leq c \left(k^{d-1} \varphi^* \left(\frac{n}{k} \right) + k^{1-\varepsilon'} n^{d-1+\varepsilon'} + k^d \varphi^{(j-1)} \left(\frac{n}{k} \right) \right)$$

(for $j = 0$ we put on the right-hand side $\varphi^{(-1)} = 0$). Moreover, we have $\varphi^*(m) = \varphi(m) + O(m^{d-1+\varepsilon'})$ and $\varphi(m) = \varphi^{(s-1)}(m)$. Thus, choosing a different threshold parameter $k = k_j$ for each j , we can rewrite the recurrence as

$$\varphi^{(j)}(n) \leq c \left(k_j^{d-1} \varphi^{(s-1)} \left(\frac{n}{k_j} \right) + k_j^{1-\varepsilon'} n^{d-1+\varepsilon'} + k_j^d \varphi^{(j-1)} \left(\frac{n}{k_j} \right) \right), \tag{2}$$

for $j = 0, \dots, s - 1$, and for a different constant $c > 0$.

We claim that the solution of this recurrence is given by

$$\varphi^{(j)}(n) \leq B_j n^{d-1+\varepsilon}, \tag{3}$$

for any $\varepsilon > 0$, where B_j is a constant depending on j , ε , d , b , and s . This is proved by holding ε fixed, by choosing $\varepsilon' = \varepsilon/2$, and by using induction on n and j . We choose

$$k_1 = k_0^\varepsilon, \quad k_2 = k_1^\varepsilon, \quad k_3 = k_2^\varepsilon, \quad \dots, \quad k_{s-1} = k_{s-2}^\varepsilon, \tag{4}$$

and choose k_0 to be sufficiently large (see below). We also define

$$B_j = (3c)^j (k_1 k_2 \cdots k_j)^{1-\varepsilon} B_0 = (3c)^j k_0^{\varepsilon - \varepsilon^{j+1}} B_0, \tag{5}$$

for $j = 1, \dots, s - 1$. To prove (3), it suffices to show, using (2) and the induction hypothesis, that, for $j = 0, \dots, s - 1$, we have

$$\frac{c B_{s-1}}{k_j^\varepsilon} + \frac{c k_j^{1-\varepsilon/2}}{n^{\varepsilon/2}} + \frac{c k_j B_{j-1}}{k_j^\varepsilon} \leq B_j,$$

or, using (4) and (5),

$$\frac{c(3c)^{s-1} k_0^{\varepsilon - \varepsilon^s} B_0}{k_0^{\varepsilon^{j+1}}} + \frac{c k_0^{\varepsilon(1-\varepsilon/2)}}{n^{\varepsilon/2}} + \frac{c k_0^\varepsilon (3c)^{j-1} k_0^{\varepsilon - \varepsilon^j} B_0}{k_0^{\varepsilon^{j+1}}} \leq (3c)^j k_0^{\varepsilon - \varepsilon^{j+1}} B_0,$$

or

$$c(3c)^{s-1}k_0^{\varepsilon-\varepsilon^s-\varrho^{s+1}} + \frac{ck_0^{\varrho(1-\varepsilon/2)}}{B_0n^{\varepsilon/2}} \leq \frac{2}{3}(3c)^jk_0^{\varepsilon-\varrho^{j+1}}.$$

Now choose k_0 so that $k_0^{\varepsilon^s} = (3c)^s$. With this choice we have, as is easily verified,

$$c(3c)^{s-1}k_0^{\varepsilon-\varepsilon^s-\varrho^{s+1}} \leq \frac{1}{3}(3c)^jk_0^{\varepsilon-\varrho^{j+1}},$$

so it remains to ensure that

$$\frac{ck_0^{\varrho(1-\varepsilon/2)}}{B_0n^{\varepsilon/2}} \leq \frac{1}{3}(3c)^jk_0^{\varepsilon-\varrho^{j+1}}. \tag{6}$$

When $j = 0$, the left-hand side of (6) assumes its maximum value and the right-hand side its minimum value, so it suffices to ensure that

$$\frac{ck_0^{1-\varepsilon/2}}{B_0n^{\varepsilon/2}} \leq \frac{1}{3},$$

which will hold if we choose, e.g., $B_0 \geq 3ck_0^{1-\varepsilon/2}$. Hence, an appropriate choice of parameters guarantees the validity of the induction step, and thus completes the proof of (3). This also completes the proof of theorem. \square

Remark. The constants B_j grow fast with s . A closer inspection of our analysis shows that they grow roughly as c^{s/ε^s} , for some constant c (which is yielded by the induction hypothesis for $d' < d$). The dependence of these constants on d and on the maximum degree of the surfaces and of their boundaries (ignoring the effect of s) is much milder.

3. Efficient Construction of the Envelope in Three Dimensions

In this section we consider the task of constructing the lower envelope of a collection Σ of n surface patches in 3-space that satisfy conditions (i)–(iv). In more precise terms, our goal is to construct the planar minimization diagram \mathcal{M}_Σ , and to label each of its faces, edges, and vertices with the appropriate combination of surfaces and/or surface boundaries that define that feature of \mathcal{M} . We do this by applying the following rather standard randomized technique. Fix some sufficiently large constant $r < n$, and draw a random sample \mathcal{R} of r surfaces from Σ . Compute (say, by any brute-force method) the lower envelope $E_{\mathcal{R}}$, and construct the *vertical decomposition* $\mathcal{V}(\mathcal{M}_{\mathcal{R}})$ of the resulting minimization diagram $\mathcal{M}_{\mathcal{R}}$;

that is, draw a vertical line segment (parallel to the y -axis) through each vertex of $\mathcal{M}_{\mathcal{R}}$, and through each locally x -extremal point on any edge of $\mathcal{M}_{\mathcal{R}}$, and extend it up and down (in the y -direction) until it meets another arc of $\mathcal{M}_{\mathcal{R}}$ (or, failing that, all the way to $\pm\infty$). This produces a decomposition of the plane into cells of “constant description complexity”: each cell is bounded from above and from below (in the y -direction) by portions of two arcs of $\mathcal{M}_{\mathcal{R}}$, and on its left and on its right by two segments parallel to the y -axis (some of these features may be missing in some of the cells). The number of these cells is proportional to the combinatorial complexity of $\mathcal{M}_{\mathcal{R}}$, namely, it is $O(r^{2+\epsilon})$, for any $\epsilon > 0$. Finally, for each cell c of $\mathcal{V}(\mathcal{M}_{\mathcal{R}})$ construct a three-dimensional cell c^* consisting of all points whose xy -projections fall in c and which lie below $E_{\mathcal{R}}$. By construction, c^* is a semiunbounded vertical cylindrical-like region, bounded from above by a portion of the single surface of Σ which attains $E_{\mathcal{R}}$ over c (or, if no such surface exists, c^* is the full cylinder consisting of all points whose xy -projections fall in c). Denote by $\mathcal{M}_{\mathcal{R}}^*$ the resulting collection of these cylindrical cells c^* .

Since no cell of $\mathcal{M}_{\mathcal{R}}^*$ intersects any surface of \mathcal{R} in its interior, and since each such cell is defined in terms of only a constant number of surfaces of Σ , it follows, using standard probabilistic arguments [10], [23], that, with high probability, no cell of $\mathcal{M}_{\mathcal{R}}^*$ intersects in its interior more than $O((n/r) \log r)$ surfaces of Σ . We now compute, for each cell c^* of $\mathcal{M}_{\mathcal{R}}^*$, the set of surfaces of Σ that meet the interior of c^* ; since r is assumed to be constant, this can be done using a brute-force method, in overall $O(n)$ time. If any of these sets exceed the bound $O((n/r) \log r)$, we simply discard \mathcal{R} and start afresh with a new random sample. In a constant expected number of trials, we will have obtained a “good” sample \mathcal{R} . We then compute recursively, at each cell c^* , the envelope E_{Σ} over its base c , and then “glue” together all these pieces of the envelope to obtain the overall envelope E_{Σ} .

If we denote by $T(n)$ the expected time needed by the algorithms to compute the lower envelope of n surface patches satisfying conditions (i)–(iv) (for some fixed b and s), we obtain the following recurrence,

$$T(n) = O(n) + O(r^{2+\epsilon}) \cdot T\left(O\left(\frac{n}{r} \log r\right)\right),$$

whose solution is easily seen to be $T(n) = O(n^{2+\epsilon})$. (We are assuming here a model of computation in which primitive operations involving a constant number of the given surfaces can be performed in constant time. Such operations include: testing whether a point lies above, on, or below a given surface, computing the intersections of a triple of surfaces, computing the intersection curve of a pair of surfaces, computing the intersection curve between one surface and the vertical-wall erected from the boundary of another surface, and so on. Since we have assumed that all the surfaces and their boundaries are algebraic of constant degree, such a model is plausible, using standard machinery from real algebraic geometry for performing each operation of this kind in an exact manner and in constant time.)

We have thus shown:

Theorem 3.1. *The lower envelope of n surface patches in three dimensions, satisfying conditions (i)–(iv) above, can be computed in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, in an appropriate “algebraic” model of computation.*

Remark. The algorithm can be made deterministic, without affecting its asymptotic running time, by using the deterministic technique or Matoušek [24] for constructing ε -nets of constant size in linear time. We have preferred the randomized approach since it is much simpler in practice.

Remark. In attempting to extend the above technique to higher dimensions, we face a technical difficulty that involves the vertical decomposition of the minimization diagram $\mathcal{M}_{\mathcal{R}}$ of a sample \mathcal{R} of surfaces: whereas in two dimensions the complexity of such a decomposition is larger than the complexity of the undecomposed diagram by only a constant factor, no similar bounds are known in higher dimensions (see [7] for the best-known bounds on such decompositions). Once this bottleneck is resolved satisfactorily, the rest of the above technique can easily be extended to higher dimensions. We also note that recently, after the original preparation of this paper, Agarwal *et al.* [1] have obtained a randomized algorithm for computing lower envelopes in four dimensions, whose expected running time is $O(n^{3+\varepsilon})$ for any $\varepsilon > 0$. They also obtain a randomized algorithm that computes all zero-, one-, and two-dimensional faces of the lower envelope, in any dimension d , with expected running time $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$. Several other related results have also appeared recently: Boissonnat and Dobrindt [6] give a randomized incremental algorithm for constructing lower envelopes in 3-space, whose expected running time is $O(n^{2+\varepsilon})$. Another such algorithm is given by de Berg *et al.* [14]. A very simple deterministic divide-and-conquer algorithm for constructing lower envelopes in 3-space is given by Agarwal *et al.* [2]; it also runs in time $O(n^{2+\varepsilon})$.

4. Applications

In this section we consider several applications of the new bounds derived in this paper.

k-Levels in Arrangements

Our results can be applied to extend the analysis of [29]:

Theorem 4.1. *Let Σ be a collection of n $(d - 1)$ -dimensional algebraic surface patches in d -space, which satisfy conditions (i)–(iii) of Section 2. Then the number of vertices of $\mathcal{A}(\Sigma)$ that lie at level at most k is $O(k^{1-\varepsilon}n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. The analysis of [29], similar to that used in this paper, implies that the number in question is $O(k^d \varphi^*(n/k))$, where $\varphi^*(r)$ is, as defined in Section 2, the

maximum number of vertices on the lower envelope of r surfaces satisfying the same conditions as Σ . The assertion is thus an immediate consequence of Theorem 2.2. □

Voronoi Diagrams

As observed in [15], Voronoi diagrams correspond to lower envelopes in arrangements in a rather natural way. In fairly full generality, we are given a set $S = \{s_1, \dots, s_n\}$ of n objects in d -space, where we assume that each object in S has “constant description complexity,” meaning, as above, that it is a semialgebraic set which is defined by a constant number of polynomials of constant degree. We are also given a metric ρ which is an algebraic function of constant degree. The Voronoi diagram of S under the metric ρ is a partition of d -space into maximal connected cells of various dimensions, so that for each cell c a subset $S' \subseteq S$ of objects exists so that for each $x \in c$ we have

$$S' = \{s' \in S: \rho(x, s') = \min\{\rho(x, s): s \in S\}\},$$

and c is a maximal connected (relatively open) cell with this property.

Following the observation of [15], it is clear that the Voronoi diagram is the minimization diagram of the collection of surfaces $\sigma_i: x_{d+1} = \rho(x, s_i)$, for $i = 1, \dots, n$. The above assumptions imply that these surfaces satisfy conditions (i)–(iii) of Section 2 (note that these are full surfaces, not surface patches). Thus our analysis yields the following results:

Theorem 4.2. *The combinatorial complexity of the Voronoi diagram of S , in the setup just defined, is $O(n^{d+\epsilon})$, for any $\epsilon > 0$.*

As an interesting special case, we can obtain, e.g., the following result, which appears to be new:

Corollary 4.3. *The combinatorial complexity of the euclidean Voronoi diagram of n lines in 3-space is $O(n^{3+\epsilon})$, for any $\epsilon > 0$.*

This result is quite general; in fact, it can be generalized even further. For example, it applies to *dynamic* Voronoi diagrams, extending the recent analysis technique of [16] and [17] (we omit the straightforward proof):

Corollary 4.4. *Let S be a set of objects in d -space as defined above. Suppose that each $s \in S$ moves (translates and rotates) with time along a given trajectory, so that, for each $s \in S$, the function $\rho(x, s(t))$ is an algebraic function, of constant degree, in the $d + 1$ variables x and t . Then the maximum number of combinatorial changes that the Voronoi diagram of S undergoes with time is $O(n^{d+1+\epsilon})$ for any $\epsilon > 0$.*

We remark that we do not know whether the above bounds are close to being tight in the worst case. For example, the (static) Voronoi diagram of n points in

d -space under the euclidean metric corresponds to the lower envelope of n hyperplanes in $(d + 1)$ -space, and thus its complexity is only $O(n^{\lceil d/2 \rceil})$. We are not aware of any larger lower bounds in the more general setup allowed here. In particular, it is an open problem whether the complexity of generalized Voronoi diagrams in 3-space can be much more than quadratic (the upper bounds derived above are only nearly cubic). Similarly, it is open whether the number of combinatorial changes of planar dynamic Voronoi diagrams, even for points moving at constant speed along straight trajectories, can be much more than quadratic.

Lines in 3-Space

Many problems involving lines in 3-space can be reduced to lower envelopes of certain arrangements in 4-space. As an illustration, we give the following generalization of a recent result of Pellegrini [27] (see also [9]). First a bit of terminology: Let S be a set of n objects in 3-space, each having constant description complexity (in the sense defined above). We say that a line l in 3-space passes above an object s if the relatively open vertical half-plane bounded from below by l does not intersect s . We can parametrize lines as points in projective 4-space, so that the fourth coordinate measures the parallel translation of the line from some canonical placement in the positive z -direction. It follows that the space of all lines that pass above all the objects of S corresponds to the region of 4-space that lies above the upper envelope of n surface patches, each representing the locus of all lines that touch an object of S (but otherwise pass above it). The combinatorial complexity of this space is defined as the complexity of this upper envelope. Thus the results of this paper are easily seen to imply:

Theorem 4.5. *The combinatorial complexity of the space of all lines that pass above all objects of S , in the setup just defined, is $O(n^{3+\epsilon})$, for any $\epsilon > 0$.*

There are several potential algorithmic applications of this bound; they all require an efficient algorithm (of near-cubic complexity) for constructing this upper envelope. As remarked above, the recent results of [1] do provide such an efficient algorithm, and consequently lead to improved solutions of several problems involving line in space, such as computing the width of a point set in 3-space (see [1] and [8] for details).

5. Conclusion

This paper and the companion paper [18] present the first general nontrivial upper bounds for the complexity of lower envelopes of algebraic surfaces or surface patches in three and higher dimensions; our bounds almost match the known lower bounds for the maximum complexity of such envelopes: they are only within a factor of $O(n^\epsilon)$ larger than the lower bounds. Thus our results almost establish the conjecture mentioned in the Introduction.

Our results raise some open questions. First, it would be nice to close, or at least to narrow, the small remaining gap between our upper bounds and the conjectured bounds. Another open problem is to extend the results given here to the problem of bounding the complexity of a single cell in an arrangement of algebraic surfaces or surface patches. Successful attempts of this kind, in the three-dimensional case, were recently made in [19] and [20], yielding a near-quadratic bound on the complexity of such a cell, and a near-quadratic algorithm for computing a single cell in arrangements of certain special classes of surfaces.

It would also be useful to restate conditions (i)–(iv) of Section 2, so that they become purely topological, and thus extend our results to apply to collections of more general surfaces.

Another problem is to extend the algorithmic technique of Section 3 to obtain efficient algorithms for computing lower envelopes in $d > 3$ dimensions (overcoming the technical difficulty mentioned at the end of Section 3). We have already noted there that some progress on this problem has recently been made by Agarwal *et al.* [1].

Finally, the technique used in this paper is relatively simple and is based essentially on a single new idea (namely, a specific application of the random sampling technique for deriving bounds on generalized “($\leq k$)-sets” in arrangements), which appears to be fairly general. A grand challenge is to investigate the applicability of this technique to other problems concerning complexity in arrangements, such as obtaining bounds for the complexity of a level, or of many cells, or of vertical decompositions of arrangements, and to related problems, such as the complexity of planar dynamic Voronoi diagrams and of generalized Voronoi diagrams in three dimensions (attempting to improve the bounds stated in Section 4 by roughly an order of magnitude).

In closing we note that, after the original preparation of this paper, more progress has been made in applying the technique and the results of this paper to a variety of related problems. These new developments include the already-mentioned papers [1], [19], [20], and also:

- (i) An application of our bounds to the analysis of the number of topologically different orthographic or perspective views of a polyhedral terrain [3].
- (ii) An efficient technique for ray shooting amidst spheres in three dimensions [25].
- (iii) Near-quadratic bounds on the complexity of the overlay of minimization diagrams of two collections of bivariate functions, and its applications to bounding the complexity of the space of common plane transversals of a collection of sets in 3-space [2].

Acknowledgments

The author wishes to express his thanks to Dan Halperin for useful discussions on this problem, and for our joint work [18] that contains the initial ideas that have led to the results of the present paper.

References

1. P. K. Agarwal, B. Aronov, and M. Sharir, Computing envelopes in four dimensions with applications, *Proc. 10th ACM Symp. on Computational Geometry*, 1994, pp. 348–358.
2. P. K. Agarwal, O. Schwarzkopf, and M. Sharir, The overlay of lower envelopes and its applications, Manuscript, 1993.
3. P. K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Proc. 5th Canadian Conference on Computational Geometry*, 1993, pp. 55–61. (To appear in *Discrete Comput. Geom.*)
4. P. K. Agarwal, M. Sharir, and P. Shor, Sharp upper and lower bounds for the length of general Davenport–Schinzel sequences, *J. Combin. Theory Ser. A* **52** (1989), 228–274.
5. J. Bochnak, M. Coste, and M-F. Roy, *Géométrie Algébrique Réelle*, Springer-Verlag, Berlin, 1987.
6. J. D. Boissonnat and K. Dobrindt, On-line randomized construction of the upper envelope of triangles and surface patches in \mathbb{R}^3 , Manuscript, 1993.
7. B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Proc. 16th Internat. Colloq. on Automata, Languages and Programming*, 1989, pp. 179–193.
8. B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, Diameter, width, closest line pair, and parametric searching, *Discrete Comput. Geom.* **10** (1993), 183–196.
9. B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir, and J. Stolfi, Lines in space: combinatorics and algorithms, *Algorithmica*, to appear.
10. K. Clarkson, New Applications of random sampling in computational geometry, *Discrete Comput. Geom.* **2** (1987), 195–222.
11. K. Clarkson, H. Edelsbrunner, L. Guibas, M. Sharir, and E. Welzl, Combinatorial complexity bounds for arrangements of curves and spheres, *Discrete Comput. Geom.* **5** (1990), 99–160.
12. K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.* **4** (1989), 387–421.
13. G. E. Collins, Quantifier elimination for real closed fields by cylindrical algebraic decomposition, *Proc. 2nd GI Conf. on Automata Theory and Formal Languages*, Springer-Verlag, Berlin, 1975, pp. 134–183.
14. M. de Berg, K. Dobrindt, and O. Schwarzkopf, On lazy randomized incremental construction, *Proc. 26th ACM Symp. on Theory of Computing*, 1994, pp. 105–114.
15. H. Edelsbrunner and R. Seidel, Voronoi diagrams and arrangements, *Discrete Comput. Geom.* **1** (1986), 25–44.
16. J.-J. Fu and R. C. T. Lee, Voronoi diagrams of moving points in the plane, *Internat. J. Comput. Geom. Appl.* **1** (1991), 23–32.
17. L. Guibas, J. Mitchell, and T. Roos, Voronoi diagrams of moving points in the plane, *Proc. 17th Internat. Workshop on Graph-Theoret. Concepts in Computer Science*, Lecture Notes in Computer Science, Vol. 570, Springer-Verlag, Berlin, 1991, pp. 113–125.
18. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, this issue, pp. 313–326.
19. D. Halperin and M. Sharir, Near-quadratic bounds for the motion planning problem for a polygon in a polygonal environment, *Proc. 34th IEEE Symp. on Foundations of Computer Science*, 1993, pp. 382–391.
20. D. Halperin and M. Sharir, Almost tight upper bounds for the single cell and zone problems in three dimensions, *Proc. 10th ACM Symp. on Computational Geometry*, 1994, pp. 11–20.
21. S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes. *Combinatorica* **6** (1986), 151–177.
22. R. Hartshorne, *Algebraic Geometry*, Springer-Verlag, New York, 1977.
23. D. Haussler and E. Welzl, ϵ -nets and simplex range queries, *Discrete Comput. Geom.* **2** (1987), 127–151.
24. J. Matoušek, Approximations and optimal geometric divide-and-conquer, *Proc. 23rd ACM Symp. on Theory of Computing*, 1991, pp. 506–511.
25. S. Mohabian and M. Sharir, Ray shooting amidst spheres in three dimensions and related problems, manuscript, 1993.

26. 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.
27. M. Pellegrini, On lines missing polyhedral sets in 3-space, *Proc. 9th ACM Symp. on Computational Geometry*, 1993, pp. 19–28.
28. J. T. Schwartz and M. Sharir, On the two-dimensional Davenport–Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371–393.
29. M. Sharir, On k -sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* **6** (1991), 593–613.
30. M. Sharir and P. K. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge, to appear.

Received February 7, 1993, and in revised form December 13, 1993.