

New Bounds for Lower Envelopes in Three Dimensions, with Applications to Visibility in Terrains*

D. Halperin¹† and M. Sharir^{1,2}

¹ School of Mathematical Sciences, Tel Aviv University,
Ramat Aviv, 69978 Tel Aviv, Israel
danha@flamingo.stanford.edu
sharir@math.tau.ac.il

² Courant Institute of Mathematical Sciences, New York University,
251 Mercer Street, New York, NY 10012, USA

Abstract. We consider the problem of bounding the complexity of the lower envelope of n surface patches in 3-space, all algebraic of constant maximum degree, and bounded by algebraic arcs of constant maximum degree, with the additional property that the interiors of any triple of these surfaces intersect in at most two points. We show that the number of vertices on the lower envelope of n such surface patches is $O(n^2 \cdot 2^{c\sqrt{\log n}})$, for some constant c depending on the shape and degree of the surface patches. We apply this result to obtain an upper bound on the combinatorial complexity of the “lower envelope” of the space of all rays in 3-space that lie above a given polyhedral terrain K with n edges. This envelope consists of all rays that touch the terrain (but otherwise lie above it). We show that the combinatorial complexity of this ray-envelope is $O(n^3 \cdot 2^{c\sqrt{\log n}})$ for some constant c ; in particular, there are at most that many rays that pass above the terrain and touch it in four edges. This bound, combined with the analysis of de Berg *et al.* [4], gives an upper bound (which is almost tight in the worst case) on the number of topologically different orthographic views of such a terrain.

* Work on this paper by the first author has been supported by a Rothschild Postdoctoral Fellowship. Work on this paper by the second author 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.

† Current address: Robotics Laboratory, Computer Science Department, Stanford University, Stanford, CA 94305, USA.

1. Introduction

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of n surface patches in 3-space. This paper addresses the problem of bounding the combinatorial complexity of the lower envelope (pointwise minimum) of the surface patches in Σ . Assume for simplicity that each σ_i is the graph of a partially defined function $z = f_i(x, y)$, and that these functions, as well as the curves delimiting the patches, are all algebraic of constant maximum degree, and that the given surfaces are in general position (see Section 2 for precise definitions). The lower envelope, when projected onto the xy -plane, generates a planar map \mathcal{M} , called the *minimization diagram* of Σ [15], with the property that over each face of \mathcal{M} the envelope is attained by a single patch (or by no patch at all), over each edge the envelope is attained by two patches simultaneously or by the boundary of a single patch, and over each vertex of \mathcal{M} the envelope is attained by three patches simultaneously, or by the intersection of the boundary of one patch with another patch, or by a point on the boundary of one patch which lies directly below the boundary of another patch, or below an intersection curve of two other patches (so that this higher point is vertically visible from the point on the lower boundary), or by a vertex of a patch boundary (a point where two arcs forming this boundary meet). The combinatorial complexity of the envelope is defined simply as the overall number of faces, edges, and vertices of \mathcal{M} , and is denoted by $\psi(\Sigma)$.

Under the assumptions made above, it is easy to show that $\psi(\Sigma) = O(n^3)$ (with a constant of proportionality that depends 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 only about quadratic in n .

The conjecture is motivated by the fact that in two dimensions, in the case of the lower envelope of n partially defined univariate functions, sharp bounds are known for the complexity of the envelope, measured simply in terms of the number of *breakpoints* along the envelope. If each pair of the functions intersect in at most s points, then 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 [2] and [11] for more details), and is only slightly superlinear in n for any fixed s . The conjecture in three dimensions attempts to extend this bound, and asserts that the complexity of the envelope is $O(n\lambda_q(n))$, for some constant q depending on the degree and shape of the given patches. The conjecture appears to be extremely difficult, and has been proven for families of only a few types of surface patches, such as triangles, and a few other types (see [13] and [15]). Better bounds are known for the special cases of planes and balls. The problem in general has been wide open; in fact, no general bounds better than $O(n^3)$ were known so far.

In this paper we obtain a subcubic bound for the complexity $\psi(\Sigma)$, provided the given surface patches are such that the interiors of any three of them intersect in at most *two* points. In fact, our bound is close to quadratic, so we almost

establish the conjecture in this special case. This property holds in several applications. As a matter of fact, if the given surfaces had this property and were full surfaces without boundary, then the results of [15] would imply (with a few additional mild assumptions) that the complexity of their envelope is $O(n^2)$. However, the fact that they are surface patches makes the analysis more difficult, and this bound is not known in that case.

Our main result is that, under the above assumptions, the complexity of the lower envelope of these surfaces is $O(n^2 \cdot 2^{c\sqrt{\log n}})$, for some constant c that depends on the degree and shape of the given surfaces. The proof is not difficult, and relies on the randomized technique of [7] and [16] for obtaining generalized “($\leq k$)-set” bounds in arrangements. This result still leaves a small gap from the conjectured complexity, but is nevertheless a significant and rather decisive step toward the establishment of the conjecture.

In a companion paper [17] the technique presented in this paper is extended to obtain similar, almost-tight bounds in more general setups, both when the maximum number of points of intersection between a triple of surfaces can be larger than two (but still remains a constant), and when the dimension is larger than three.

Our result was motivated by, and has an interesting application to, the problem mentioned in the abstract, namely, the problem of bounding the combinatorial complexity of the “lower envelope” of the space of all rays in 3-space that lie above a given polyhedral terrain K with n edges. This envelope consists of all rays that touch the terrain (but otherwise lie above it). We show that the combinatorial complexity of the envelope is $O(n^3 \cdot 2^{c\sqrt{\log n}})$ for some constant c ; in particular, there are at most that many rays that pass above the terrain and touch it in four edges. This bound is derived by applying the three-dimensional envelope result n times, each time fixing an edge e of K and considering the three-dimensional space of all lines (or rays) that touch e but otherwise pass above K . It is fairly easy to show (and indeed is shown in Section 3) that the conditions required by our analysis hold in this application. Our bound for the complexity of the space of lines passing above a terrain was recently, and independently, obtained by Pellegrini [14], but it is not clear whether his result can be extended to the space of rays passing above a terrain.

This bound on the complexity of the envelope of the space of rays over a terrain has an interesting application to the problem of bounding the number of topologically different orthographic views of such a terrain. A recent study by de Berg *et al.* [4] gives a bound of $O(n^5 \cdot 2^{\alpha(n)})$ for this number, but unfortunately their argument turned out to be erroneous. A careful (and correct) restatement of the result of de Berg *et al.* is that the bound on the number of topologically different orthographic views of a terrain is $O(n\lambda_4(n)\mu(n))$, where $\mu(n)$ is the maximum complexity of the ray-envelope of a terrain with n edges. Plugging in our bound for $\mu(n)$, we get the bound $O(n^5 \cdot 2^{c\sqrt{\log n}})$ for the number of views, for some absolute constant c . Except for a remaining small factor, this result almost matches the upper bound originally asserted in [4], and, as shown in [4], is almost tight in the worst case.

2. Lower Envelopes in 3-Space

Let $\Sigma = \{\sigma_1, \dots, \sigma_n\}$ be a given collection of n surface patches in 3-space that satisfy the following conditions:

- (i) Each σ_i is monotone in the xy -direction (that is, every vertical line intersects σ_i in at most one point). Moreover, each σ_i is a portion of an algebraic surface of constant maximum degree b .
- (ii) The vertical projection of σ_i onto the xy -plane is a planar region bounded by a constant number of algebraic arcs of constant maximum degree (say, b too).
- (iii) The relative interiors of any triple of the given surfaces intersect in at most two points.
- (iv) The surface patches in Σ are in *general position*. This excludes degenerate configurations where four surfaces meet at a point, the boundaries of two surfaces meet, the boundary of one surface meets an intersection curve of two other surfaces, a pair of surfaces are tangent to each other, a singular point on one surface lies on the boundary of another surface or on an intersection curve between two other surfaces, etc.

Conditions (i)–(iii) are essential for the analysis, while condition (iv) is made to simplify the forthcoming arguments. It involves no real loss of generality, because, as can be shown (see the companion paper [17] for details), the maximum complexity of the envelope is achieved, up to a constant factor, when the given surfaces are in general position. We also note that the first part of condition (i) is not essential, because we can always cut a surface σ along the constant number of arcs which are the loci of points on σ having vertical tangency, to obtain a constant number of xy -monotone surfaces whose union is σ .

The *lower envelope* of Σ is the graph of the (partial) function $z = E_\Sigma(x, y)$ that maps each point (x, y) to the height of the lowest point of intersection between the vertical line through (x, y) and the surfaces in Σ (if that line meets no surface, the function is undefined at (x, y)). If we project E onto the xy -plane we obtain a planar map, denoted by $\mathcal{M} = \mathcal{M}_\Sigma$, having the properties stated in the Introduction.

Theorem 2.1. *The combinatorial complexity $\psi(\Sigma)$ of the lower envelope of a collection Σ of n surface patches that satisfy conditions (i)–(iv) is*

$$O(n^2 \cdot 2^{c\sqrt{\log n}}),$$

for some constant c that depends on the degree b and the shape of the given surfaces.

Proof. Let us denote by $\psi(n)$ the maximum value of $\psi(\Sigma)$, taken over all collections Σ of n surface patches that satisfy conditions (i)–(iv) (for a fixed degree b).

We distinguish between *inner vertices* of the envelope, namely, vertices of the envelope that are intersection points of the relative interiors of three surface patches in Σ , and the remaining *outer vertices*. Consider first the outer vertices.

Lemma 2.2. *The maximum number of outer vertices is $O(n\lambda_{q+2}(n))$, for some constant q depending on the degree and shape of the given surfaces.*

Proof. Let us assume general position. There are $O(n)$ original vertices of the given patches, and $O(n^2)$ points of intersection between the boundary of one surface patch and the relative interior of another patch. Any other outer vertex is formed either when the boundary of one surface patch passes above the boundary of another patch, or when an intersection curve of two patches passes above the boundary of a third patch. The number of outer vertices of the first kind is clearly $O(n^2)$. As to outer vertices of the second kind, fix a surface patch $\sigma \in \Sigma$. We claim that the total number of such vertices that lie on the boundary of the fixed σ is $O(\lambda_{q+2}(n))$, for some constant q depending on the degree and shape of the surfaces. This is shown as in [5] and [6]: Let α be one of the (constant number of) algebraic arcs that form the boundary of σ , and let H be the vertical surface formed by the union of all vertical rays whose bottom endpoints lie on α . For each surface $\sigma_i \in \Sigma - \{\sigma\}$, let $\delta_i = \sigma_i \cap H$. The properties that the surfaces satisfy imply that each δ_i is the union of a constant number of connected arcs, and that each pair of such arcs intersect in at most some constant number, q , of points. It is easily checked that each of the vertices under consideration must arise as a *breakpoint* in the lower envelope of the arcs δ_i , over one of the boundary pieces α of σ . Hence the total number of such endpoints is $O(\lambda_{q+2}(n))$ [11], as asserted. Summing over all $\sigma \in \Sigma$, the total number of outer vertices of the second kind is $O(n\lambda_{q+2}(n))$. □

Hence the number of outer vertices of all kinds is nearly quadratic and smaller than the bound asserted in the theorem. The number of edges of \mathcal{M}_Σ that do not contain any vertex at all is easily seen to be $O(n^2)$. Thus, if we establish the bound in the theorem for the number of inner vertices, and apply Euler's formula for planar maps, we easily conclude that the same bound also holds for the total complexity of \mathcal{M}_Σ .

Let p be an inner vertex of the lower envelope, formed by the intersection of (the relative interiors of) three surface patches $\sigma_1, \sigma_2, \sigma_3$. By condition (iii), these patches intersect in at most one more point, so, with no loss of generality, we can assume that there is no such intersection in the half-space $x > x(p)$, where $x(p)$ is the x -coordinate of p . By the general position assumption, we may assume that p is not a singular point of any of these patches, and that they meet each other transversally at p . Then, locally near p , the lower envelope is approximated by the lower envelope of the three tangent planes to $\sigma_1, \sigma_2, \sigma_3$ at p . In particular, each of the intersection curves $\gamma_{12} = \sigma_1 \cap \sigma_2, \gamma_{13} = \sigma_1 \cap \sigma_3, \gamma_{23} = \sigma_2 \cap \sigma_3$ contains a maximal connected x -monotone arc that emanates from p and is hidden from below by the third surface. Let $\beta_{12}, \beta_{13}, \beta_{23}$ denote these

respective arcs, and let z_{12}, z_{13}, z_{23} denote the other endpoints of these arcs. It also follows that the positive span of the xy -projections of the three outgoing tangent directions of $\beta_{12}, \beta_{13}, \beta_{23}$ is the entire xy -plane: indeed, the three vectors oppositely oriented to these directions lie along the edges emerging from p of the lower envelope of the three tangent planes to the surfaces at p , and the xy -projections of these vectors cannot lie in a single half-plane, as is easily checked. Hence, at least one of these arcs, say β_{12} , emanates from p in the positive x -direction.

Two cases can arise:

- (a) z_{12} is an endpoint of (some connected x -monotone portion of) γ_{12} . In this case we charge p to z_{12} . Our assumptions imply that each γ_{12} has only a constant number of such endpoints, so the total number of such charges will be proportional to the total number of intersection curves, namely, $O(n^2)$.
- (b) z_{12} is a point along γ_{12} that lies directly above the boundary of σ_3 (see Fig. 1). The difficulty is that z_{12} need not be vertically visible from $\partial\sigma_3$ —there may be many other surfaces lying below z_{12} and above the boundary (the boundary itself need not be on the lower envelope at this point, but this does not concern us in the analysis given below). Suppose there are t such surfaces. We fix some threshold parameter $k > 0$, to be determined shortly, and consider the following two subcases:
 - (i) $t > k$. 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. Our assumption implies that β_{12} contains at least t vertices of $\mathcal{A}(\Sigma^*)$, and we charge p to the block of the first k of these vertices in




Fig. 1. Case (b) of the proof: z_{12} lies above the boundary of σ_3 . The envelope is shown from below; solid edges are visible whereas dashed edges are hidden from the envelope.

their order from p to z_{12} along β_{12} . (Indeed, for any extended surface σ^* that hides z_{12} from $\partial\sigma_3$, we trace β_{12} from p toward z_{12} ; since σ^* does not lie below p but lies below z_{12} , there has to be a point along β_{12} where this curve intersects σ^* , yielding a vertex of $\mathcal{A}(\Sigma^*)$ along β_{12} .) Each vertex of $\mathcal{A}(\Sigma^*)$ can be charged in this manner only a constant number of times: By the general position assumption, each inner vertex v is the intersection point of exactly three surfaces, and hence it is the meeting point of three intersection curves (of pairs of the three surfaces). Evidently, for each of these three curves, there are at most two vertices p lying on the curve and on the lower envelope for which the vertex v will be charged in this manner.

Define the *level* of a point x of 3-space in $\mathcal{A}(\Sigma^*)$ to be the number of surfaces of Σ^* that lie strictly below x (which is the same as the number of original surfaces in Σ that lie below x). It is easily checked that each of the charged vertices along β_{12} is at level at most k (indeed, as we trace β_{12} from p to z_{12} , the level can change (by ± 1) only when we cross one of these vertices).

Our goal is thus to obtain an upper bound for the number of vertices of $\mathcal{A}(\Sigma^*)$ that lie at level $\leq k$. For this we apply the analysis technique of [7] and [16]. That is, we choose a random sample \mathcal{R} of $r = \lfloor n/k \rfloor$ surfaces of Σ^* , and construct their arrangement $\mathcal{A}(\mathcal{R})$. Let v be a vertex of $\mathcal{A}(\Sigma^*)$ at level $\xi \leq k$. The probability that v shows up as a vertex of the lower envelope $E_{\mathcal{R}}$ of the surfaces in \mathcal{R} is $\binom{n-\xi-3}{r-3} / \binom{n}{r}$: out of the total number $\binom{n}{r}$ of possible samples, those that make v appear as such a vertex are precisely those that contain the three surfaces defining v and do not contain any of the ξ surfaces lying below v . Hence, we have

$$\sum_{\xi=0}^k \frac{\binom{n-\xi-3}{r-3}}{\binom{n}{r}} F_{\xi} \leq \psi(\mathcal{R}) \leq \psi(r),$$

where F_{ξ} is the number of vertices v of $\mathcal{A}(\Sigma^*)$ at level ξ . This can be rewritten as

$$\frac{r(r-1)(r-2)}{n(n-1)(n-2)} \cdot \sum_{\xi=0}^k \frac{(n-r)(n-r-1)\cdots(n-r-\xi+1)}{(n-3)(n-4)\cdots(n-\xi-2)} F_{\xi} \leq \psi(r),$$

which implies

$$\left(\frac{n-r-k+1}{n-k-2}\right)^k \cdot \left(\sum_{\xi=0}^k F_{\xi}\right) \leq \frac{n(n-1)(n-2)}{r(r-1)(r-2)} \psi(r).$$

As in [7] and [16], it can be easily verified that for $r = n/k$ we have

$$\sum_{\xi=0}^k F_{\xi} = O\left(k^3\psi\left(\frac{n}{k}\right)\right);$$

in other words, the number of vertices of $\mathcal{A}(\Sigma^*)$ at level $\leq k$ is $O(k^3\psi(n/k))$, which in turn implies that the number of inner vertices p of E_{Σ} in this subcase is $O(k^2\psi(n/k))$.

- (ii) $t \leq k$. Let α be the algebraic arc on the boundary of σ_3 which lies below z_{12} , and let H be the vertical surface formed by the union of all vertical rays whose bottom endpoints lie on α . For each surface $\sigma_i \neq \sigma_3$, let $\delta_i = \sigma_i \cap H$. As above, the properties that the surfaces satisfy imply that each δ_i is the union of a constant number of connected arcs, and that each pair of such arcs intersect in at most q points. Clearly, the point z_{12} must be a vertex of the two-dimensional arrangement \mathcal{A}_{α} of the arcs δ_i within H . The level of z_{12} in \mathcal{A}_{α} is defined, in complete analogy to the definition in the preceding subcase, to be the number of arcs that intersect the downward-directed ray emanating from z_{12} (or, rather, the downward-directed segment from z_{12} to α). Clearly, the level of z_{12} in \mathcal{A}_{α} is at most k . We can thus apply Theorem 1.3 of [16], which is proved using a random-sampling argument similar to the one in the preceding subcase, and which asserts that the maximum number of vertices at level at most k in such an arrangement is $O(k^2\lambda_{q+2}(n/k))$. Hence the number of vertices z_{12} , over all choices of σ_1 and σ_2 but with a fixed σ_3 , is $O(k^2\lambda_{q+2}(n/k))$. We charge p to z_{12} ; since each such z_{12} can be charged in this manner at most once, it follows that the overall number of inner vertices in the present subcase is

$$O\left(nk^2\lambda_{q+2}\left(\frac{n}{k}\right)\right) = O(nk\lambda_{q+2}(n)).$$

Thus, if we add all the bounds obtained so far, including those for the number of outer vertices, we obtain the following recurrence for ψ :

$$\psi(n) \leq Ak^2\psi\left(\frac{n}{k}\right) + Ank\lambda_{q+2}(n), \tag{1}$$

for some absolute positive constant A . We claim that the solution to this recurrence is

$$\psi(n) \leq n^2 \cdot 2^{c\sqrt{\log n}}, \tag{2}$$

for a sufficiently large constant c . The proof is by induction on n . First, by choosing c sufficiently large, we can assume that (2) holds for all $n \leq n_0$, where n_0 is chosen so that $\lambda_{q+2}(n) \leq (n/2A) \cdot 2^{\sqrt{\log n}}$ for all $n > n_0$ (this is

always possible by the results of [2]). For $n > n_0$, choose $k = 2^{(c-1)\sqrt{\log n}}$. The inequality (1) and the induction hypothesis imply that

$$\psi(n) \leq Ak^2 \left(\frac{n}{k}\right)^2 \cdot 2^{c\sqrt{\log(n/k)}} + Ank\lambda_{q+2}(n),$$

so it suffices to show that

$$A \cdot 2^{c\sqrt{\log n - (c-1)\sqrt{\log n}}} + \frac{1}{2} \cdot 2^{c\sqrt{\log n}} \leq 2^{c\sqrt{\log n}},$$

or that

$$c\sqrt{\log n - (c-1)\sqrt{\log n}} \leq c\sqrt{\log n} - (1 + \log A),$$

which is easily seen to hold provided we choose $c > 1 + \sqrt{2(1 + \log A)}$. This completes the proof of the theorem. \square

3. The Envelope of Lines or Rays Over a Terrain

Let K be a *polyhedral terrain* in 3-space; that is, K is a continuous piecewise-linear surface intersecting each vertical line in exactly one point. Suppose K has n edges. A line l is said to *lie over* K if every point on l lies on or above K . Let \mathcal{L}_K denote the space of all lines that lie over K . The *lower envelope* of \mathcal{L}_K consists of those lines that touch at least one edge of K . Assuming general position of the edges of K , a line in \mathcal{L}_K (or any line for that matter) can touch at most four edges of K . Our goal is to analyze the combinatorial complexity of the lower envelope. To get a feeling of what this lower envelope is, consider the four-dimensional space \mathcal{V} of parametric representation of lines in 3-space, where each point represents one nondirected nonvertical line in 3-space. For each edge e of the terrain, the points of \mathcal{V} that represent lines in contact with e lie on a three-dimensional surface patch σ_e . The collection of all these surface patches defines a partitioning of the space \mathcal{V} into cells of various dimensions. It is evident that the points of \mathcal{V} corresponding to lines in \mathcal{L}_K occupy a single connected component C of \mathcal{V} , whose boundary is determined by portions of the surfaces σ_e . We define the complexity of the lower envelope of lines in \mathcal{L}_K to be the overall number of lower-dimensional cells on the boundary of C . To simplify matters (and with no loss of generality) we only count the number of its *vertices*, namely, those corresponding to lines that touch four distinct edges of K (and we refer to them as “vertices” of \mathcal{L}_K). We show:

Theorem 3.1. *The number of vertices of \mathcal{L}_K , as defined above, is $O(n^3 \cdot 2^{c\sqrt{\log n}})$ for some absolute positive constant c .*

Proof. We fix an edge e_0 of K , and bound the number of lines of \mathcal{L}_K that touch e_0 and three other edges of K , with the additional proviso that these three other contact points all lie on one fixed side of the vertical plane passing through e_0 . We then multiply this bound by the number n of edges to obtain a bound on the overall number of vertices of \mathcal{L}_K . We want to rephrase this problem in terms of the lower envelope of a certain collection of surface patches in 3-space, one patch for each other edge of K , to which we apply the results of the previous section.

The space \mathcal{L}_{e_0} of oriented lines that touch e_0 is three-dimensional: each such line l can be specified by a triple (t, k, ζ) , where t is the point of contact with e_0 (or, more precisely, the distance of that point from one designated endpoint of e_0), and $k = \tan \theta$, $\zeta = -\cot \varphi$, where (θ, φ) are the spherical coordinates of the direction of l , that is, θ is the orientation of the xy -projection of l , and φ is the angle between l and the positive z -axis.

For each edge $e \neq e_0$ of K let σ_e be the surface patch in \mathcal{L}_{e_0} consisting of all points (t, k, ζ) representing lines that touch e and are oriented from e_0 to e . Note that if $(t, k, \zeta) \in \sigma_e$, then $\zeta' > \zeta$ iff the line (t, k, ζ') passes below e . It thus follows that a line l in \mathcal{L}_{e_0} is a vertex of the lower envelope of \mathcal{L}_K if and only if l is a vertex of the lower envelope of the surfaces σ_e in the $tk\zeta$ -space, where the height of a point is its ζ -coordinate. Hence, it remains to show that these surfaces satisfy conditions (i)–(iv) of the previous section, and then the theorem will easily follow from Theorem 2.1.

Condition (i) requires each σ_e to be monotone in the tk -direction, which is immediate by definition; the algebraicity and the constant degree of these surfaces is also easy to verify. The vertical projection of σ_e onto the tk -plane is easily seen to be the intersection of two double wedges—it is the set dual to the set of all lines in the xy -plane that intersect the xy -projections of e_0 and e , under an appropriate (and standard) duality. Hence condition (ii) is also satisfied. (Since this projection may be disconnected, we may want to replace each σ_e by a constant number of subpatches, so that the tk -projection of each subpatch is a convex polygon of at most four sides.) Condition (iii), which is the crucial one, follows from the observation that a point of intersection of the relative interiors of three surfaces $\sigma_{e_1}, \sigma_{e_2}, \sigma_{e_3}$, corresponds to a line that passes through the four edges e_0, e_1, e_2 , and e_3 , and it is well known that there can be at most two such lines, assuming that these four edges are in general position (see, e.g., [12]). Condition (iv) can be enforced by assuming the terrain K to be in general position. We argue, as in [17], that the maximum complexity of the envelope is achieved, up to a constant factor, when the terrain, and hence the surfaces defining the lower envelope of \mathcal{L}_K , are in general position.

Hence, putting everything together and applying Theorem 2.1, we readily obtain the bound asserted in the theorem. \square

Remarks. (1) The bound of Theorem 3.1 has been independently obtained by Pellegrini [14], using a different proof technique.

(2) Recently, de Berg [3] has shown a lower-bound construction of complexity




Fig. 2. An $\Omega(n^3)$ construction for the lower envelope of \mathcal{L}_K [3].

$\Omega(n^3)$ for the envelope of \mathcal{L}_K , implying that our upper bound is almost tight in the worst case. The construction consists of an almost flat (horizontal) “hill” with $n/3$ edges, and two sets of $n/3$ steep pyramids (“spikes”) each, in front of the hill; see Fig. 2 for an illustration. Clearly, this terrain has $O(n)$ edges. If the dimensions of the hill and spikes are assigned appropriately (in particular, the edges of the hill are long enough, and the spikes are sufficiently high), then a line lying over the terrain can touch an edge of the hill and (an edge of) one spike of each row simultaneously, and this triple contact still leaves one degree of freedom for motion of the line along a small segment of each of the edges involved. This holds for every triple of an edge of the hill, and one spike from either row. Thus we have $\Omega(n^3)$ distinct one-dimensional edges on the lower envelope of \mathcal{L}_K . This construction induces $\Omega(n^3)$ vertices on the lower envelope of \mathcal{L}_K which are obtained when a line touches two horizontal edges of the hill bounding the same face, and two spikes, one from each row.

We can extend the result of Theorem 3.1 as follows. Let K be a terrain as above. Let \mathcal{R}_K denote the space of all rays in 3-space with the property that each point on such a ray lies on or above K . We define the lower envelope of \mathcal{R}_K and its vertices in complete analogy to the case of \mathcal{L}_K . By inspecting the proof of Theorem 3.1, it can be easily verified that it applies equally well to rays instead of lines. This is because, after fixing an edge e_0 , each “ray-vertex” of \mathcal{R}_K under consideration, when extended into a full line, becomes a “line-vertex” of \mathcal{L}_{K^*} , where K^* is the portion of K cut off by a half-space bounded by the vertical plane through e_0 . Hence we obtain:

Corollary 3.2. *The number of vertices of \mathcal{R}_K , as defined above, is also $O(n^3 \cdot 2^{c\sqrt{\log n}})$.*

This corollary is needed in the following subsection.

3.1. The Number of Orthographic Views of a Polyhedral Terrain

We next apply Theorem 3.1 to obtain a bound on the number of topologically different orthographic views (i.e., views from infinity) of a polyhedral terrain K with n edges. This problem has been studied by de Berg *et al.* [4]. However, there is a certain technical error in their analysis (which we explain below). As mentioned above, a correct form of the bound is $O(n\lambda_4(n)\mu(n))$, where $\mu(n)$ is the complexity of the ray-envelope of a terrain with n edges. In this subsection we explain the connection between $\mu(n)$ and the bound on the number of views, and using the result of the previous subsection, we conclude that this bound is $O(n^5 \cdot 2^{c'\sqrt{\log n}})$, for a constant c' slightly larger than the constant c in the bound of Corollary 3.2.

Following the analysis of [4], each orthographic view of K can be represented as a point on the sphere at infinity \mathcal{S}^2 . For each triple (e_1, e_2, e_3) of edges of K we consider the locus $\gamma_{(e_1, e_2, e_3)}$ of views for which these three edges appear to be concurrent (that is, a line parallel to the viewing direction which touches these three edges exists); each such locus is a curve along \mathcal{S}^2 .

We next replace each curve $\gamma = \gamma_{(e_1, e_2, e_3)}$ as above by its maximal *visible* portions; a point on γ is said to be visible if the corresponding line that touches the three edges e_1, e_2, e_3 either lies over K or else penetrates below K only at points that lie further away from its contacts with the edges e_1, e_2, e_3 ; in other words, we require the existence of a ray in the viewing direction that touches $e_1, e_2,$ and e_3 but otherwise lies fully above K . As is easily verified, each visible portion of γ is delimited either at an original endpoint of γ or at a point whose corresponding ray is a vertex of \mathcal{R}_K . Hence, by Corollary 3.2, the total number of the visible portions of the loci γ is $O(n^3 \cdot 2^{c\sqrt{\log n}})$. We refer to these visible portions as *arcs of visible triple-contact views*.

We now continue along the lines of the analysis of [4]. That is, we consider the arrangement of the arcs of visible triple-contact views, and observe that the number of views that we seek to bound is proportional to the complexity of the arrangement of these arcs within \mathcal{S}^2 . We next apply a result of Cole and Sharir [8], which, rephrased in the context under discussion, states that each meridian of \mathcal{S}^2 crosses at most $k = O(n\lambda_4(n))$ arcs. As shown in [4], this implies that the complexity of the arrangement of these arcs is $O(Nk)$, where N is the number of arcs. Hence we obtain

Corollary 3.3. *The number of topologically different orthographic views of a polyhedral terrain with n edges is*

$$O(n^4 \lambda_4(n) \cdot 2^{c'\sqrt{\log n}}) = O(n^5 \cdot 2^{c'\sqrt{\log n}}),$$

for c' slightly larger than c .

Remarks. (1) The technical error in [4] is that they considered the arrangement of the entire loci $\gamma_{(e_1, e_2, e_3)}$, whereas, for the result of [8] to apply, one has to consider, as done above, only their visible portions.

(2) de Berg *et al.* [4] also give a lower bound of $\Omega(n^5\alpha(n))$ for the number of topologically different orthographic views of a polyhedral terrain with n edges. Thus the above bound is almost tight in the worst case.

(3) The manuscript [4] also analyzes the number of topologically different *perspective* views of a polyhedral terrain. Again they make the technical error noted above, except that now they have to consider an arrangement of surface patches in 3-space rather than an arrangement of curves on the sphere at infinity. As it turns out, what is needed here is a bound on the number of vertices of the lower envelope of the space \mathcal{E}_K of all line segments that lie over K , defined in analogy with the spaces \mathcal{L}_K and \mathcal{R}_K . Unfortunately, we do not know how to extend our analysis to obtain nontrivial bounds for the case of \mathcal{E}_K . Nevertheless, Agarwal and Sharir have recently obtained an almost tight bound on the number of topologically different perspective views of a polyhedral terrain using a different approach [1].

4. Conclusion

The new bounds obtained in Section 2 for the complexity of the lower envelope of surface patches in 3-space push the frontier of research on these problems a step further. In some intuitive sense the case studied in this paper is the next simplest case for which subcubic bounds on the envelope complexity were not known. So far we do not have other applications of our result, beyond the application to terrain visibility given in Section 3, but we anticipate that such applications will be forthcoming.

There are several open problems that our study raises. The most obvious one is to close the remaining small gap between our upper bound and the known near-quadratic lower bounds for the complexity of the envelope; we continue to conjecture that the correct bound is $O(n\lambda_s(n))$, for some constant s depending on the degree and shape of the surfaces.

Another open problem is to obtain nontrivial bounds on the complexity of the space \mathcal{E}_K , as defined above.

We remark that the ideas used in the proof of Theorem 2.1 can be applied to the problem of bounding the complexity of a single cell in the free configuration space of the motion-planning problem for an arbitrary polygonal object moving (translating and rotating) in a two-dimensional polygonal environment, leading to near-quadratic bounds on that complexity. The dissertation [9] has studied several special cases of this motion-planning problem, and obtained better, often near-quadratic bounds in these cases, but no subcubic bounds were known for the general problem, as just stated. This extension of our results requires considerably more involved analysis than the one given in this paper, and it is presented in a companion paper [10].

Finally, the new technique developed in this paper is extended in the companion paper [17] to obtain similar almost tight bounds for the complexity of the envelope in cases where the maximum number of intersection points between any triple of surface patches is a constant greater than two, as well as in higher dimensions.

See also the forthcoming book [18], where all these new results, as well as many other related results and applications, are presented.

References

1. P. K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Proc. 5th Canadian Conf. on Computational Geometry*, 1993, pp. 55–60.
2. 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.
3. M. de Berg, Private communication, 1993.
4. M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld, Sparse arrangements and the number of views of polyhedral scenes, Manuscript, 1991.
5. B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Theoretical Comput. Sci.* **84** (1991), 77–105.
6. 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.
7. K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* **4** (1989), 387–421.
8. R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *J. Symbolic Comput.* **7** (1989), 11–30.
9. D. Halperin, Algorithmic Motion Planning via Arrangements of Curves and of Surfaces, Ph.D. Dissertation, Computer Science Department, Tel Aviv University, July 1992.
10. D. Halperin and M. Sharir, Near-quadratic bounds for the motion planning problem for a polygon in a polygonal environment, *Proc. 34th Ann. Symp. on Foundations of Computer Science*, 1993, pp. 382–391.
11. S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes, *Combinatorica* **6** (1986), 151–177.
12. S. L. Kleiman and D. Laksov, Schubert calculus, *Amer. Math. Monthly* **79** (1972), 1061–1082.
13. 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.
14. M. Pellegrini, On lines missing polyhedral sets in 3-space, *Proc. 9th ACM Symp. on Computational Geometry*, 1993, pp. 19–28.
15. J. T. Schwartz and M. Sharir, On the two-dimensional Davenport–Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371–393.
16. M. Sharir, On k -sets in arrangements of curves and surfaces, *Discrete Comput. Geom.* **6** (1991), 593–613.
17. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, this issue, pp. 327–345.
18. 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.