

On the Number of Views of Polyhedral Terrains*

P. K. Agarwal¹ and M. Sharir²

¹ Department of Computer Science, Duke University,
Box 90129, Durham, NC 27708-0129, USA

² School of Mathematical Sciences, Tel Aviv University,
Ramat Aviv 69 978, Israel
and

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

Abstract. We show that the number of topologically different orthographic views of a polyhedral terrain with n edges is $O(n^{5+\varepsilon})$, and that the number of topologically different perspective views of such a terrain is $O(n^{8+\varepsilon})$, for any $\varepsilon > 0$. Both bounds are almost tight in the worst case. The proofs are simple consequences of the recent almost-tight bounds of [11] on the complexity of lower envelopes in higher dimensions.

1. Introduction

Let Σ be a polyhedral terrain with n edges in 3-space. That is, Σ is the graph of a continuous bivariate piecewise-linear function. An *orthographic view* of Σ in some direction ω is defined as follows. Take a plane Π at minus infinity orthogonal to ω . For each point $q \in \Pi$, consider the ray ρ_q emanating from q in direction ω , and let ξ_q be the first point, if any, along ρ_q at which it intersects Σ . We can partition Π into a collection of maximal open connected regions, with the property that, for each region f , all points $q \in f$ are such that their “hitting points” ξ_q lie on the

* Pankaj Agarwal has been supported by National Science Foundation Grant CCR-91-06514. Micha Sharir has been supported by National Science Foundation 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.

same face of Σ (or none of these hitting points exist). The orthographic view in direction ω is the resulting planar map on Π .

Similarly, a *perspective view* of Σ is defined as follows. Take a point v in 3-space (lying above Σ), consider the collection of rays emanating from v in all possible directions, and partition the unit sphere of directions into maximal connected regions, so that all rays emanating from v whose directions lie in such a region first hit the same face of Σ (or do not hit Σ at all). The resulting spherical map is the perspective view of Σ from v .

If we vary the orthographic direction ω continuously, the corresponding view also changes continuously. However, the combinatorial and topological structure of the view remains the same until we reach a “critical direction” at which the corresponding planar map undergoes some combinatorial change (where vertices, edges, or faces newly appear or disappear); the loci of critical directions are curves of directions at which there is a ray that either passes through three distinct edges of Σ , or passes through a vertex of Σ and another edge of Σ , before piercing through Σ . These curves partition the sphere of directions into maximal connected regions, so that the topological structure of the orthographic view of Σ remains the same for all directions within each region. This partition of the sphere of directions forms a planar graph. The dual of this graph is known as the *aspect graph* of Σ . That is, there is a vertex in the aspect graph for each topologically distinct view, and two vertices v_1, v_2 are connected by an edge if the viewing direction can be varied continuously so as to pass from a view represented by v_1 to a view represented by v_2 without having any intermediate views; see [6] and [9] for a more formal definition of aspect graphs. An aspect graph for perspective views can be defined in an analogous manner. In this case it is the dual graph of a partitioning of 3-space, where the perspective views from all points in any single region are topologically equivalent. As a matter of fact, the aspect graph can be defined for any scene in 3-space. In this paper we study the size of the aspect graph (i.e., the number of topologically different views) of a polyhedral terrain. Aspect graphs have applications in computer vision, computer graphics, and object recognition [5], [6], [9].

The notion of aspect graphs was introduced by Koenderink and van Doorn [8], and also by Chakravarty and Freeman [3]. Plantinga and Dyer [9] proved that the maximum number of topologically different orthographic and perspective views of a convex polyhedron with n faces are $\Theta(n^2)$ and $\Theta(n^3)$, respectively. For a nonconvex polyhedron, and for arbitrary polyhedral scenes, the bounds are $\Theta(n^6)$ and $\Theta(n^9)$, respectively [10]. Snoeyink [12] has shown that even if we restrict the objects to be axis-parallel, the bound on the maximum number of orthographic views remains $\Theta(n^6)$.

The number of topologically different views of a polyhedral terrain was first considered by Cole and Sharir [4]. They showed that if the viewpoint moves along a vertical line, the maximum number of topologically different perspective views is $O(n^2 \cdot 2^{\alpha(n)})$, where $\alpha(n)$ is the inverse Ackermann function. They also showed that the number of topologically different views along a vertical line can be as large as $\Omega(n^2)$ in the worst case. Based on the result of Cole and Sharir, de Berg *et al.* [1] proved that the maximum number of orthographic views of a

polyhedral terrain with n faces is $O(n^5 \cdot 2^{\alpha(n)})$, and the maximum number of perspective views is $O(n^8 \cdot 2^{\alpha(n)} \log^2 n)$. However, their proof turned out to be erroneous. Recently Halperin and Sharir [7] gave an upper bound of $O(n^5 2^{c\sqrt{\log n}})$ on the number of orthographic views of a polyhedral terrain, based on their result on the complexity of the lower envelope of a set of surface patches in 3-space. Their proof does not extend to perspective views, and, moreover, it is somewhat complicated.

In this paper we give a very simple and direct proof for an upper bound of $O(n^{5+\varepsilon})$, for any $\varepsilon > 0$,¹ on the maximum number of orthographic views. It uses the recent result of Sharir [11] on the complexity of the lower envelope of surface patches in higher dimensions. Following a similar approach, we prove that the maximum number of perspective views of a polyhedral terrain is $O(n^{8+\varepsilon})$. The best-known lower bounds for the maximum number of orthographic and perspective views are $\Omega(n^2 \alpha(n))$ and $\Omega(n^8 \alpha(n))$, respectively [1], so our bounds are close to optimal. Our techniques also extend to more general terrains, where each face is a portion of an algebraic surface of some (small) fixed degree, e.g., each face is a portion of a sphere.

2. The Number of Orthographic Views

Theorem 1. *The number of topologically different orthographic views of a polyhedral terrain Σ with n edges is $O(n^{5+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. We consider the space of all spatial orientations, represented by points on the unit sphere S about the origin. Each spatial orientation $\omega \in S$ is parametrized by its spherical coordinates (θ, φ) , where θ is the horizontal orientation of the xy -projection of ω , and φ is the *azimuth* of ω , that is, the angle between the positive z -axis and the ray from O to ω . We wish to partition S into maximal connected regions, so that the topological structure of the view of Σ in all directions within each region remains the same. Let $M(\Sigma)$ denote the resulting subdivision of S . It is easily checked that the topological structure of the view of Σ changes at a direction ω if there is a ray ρ in direction ω coming from minus infinity and satisfying one of the following two conditions:

1. ρ passes through a vertex and an edge of Σ before penetrating into the open region lying below Σ (if at all).
2. ρ passes through three edges of Σ before penetrating into the open region lying below Σ (if at all).

Therefore, for each direction ω lying on an edge of $M(\Sigma)$, there is a ray in direction ω that satisfies one of these two conditions. Moreover, a vertex of $M(\Sigma)$ corresponds to a direction ω such that there are two parallel rays coming from minus

¹ Throughout this paper, ε denotes a positive constant which can be chosen arbitrarily small with an appropriate choice of the constant of proportionality.

infinity in direction ω , each of which satisfies one of the above conditions. To bound the number of different topological views of Σ , it suffices to obtain a bound on the number of vertices in $M(\Sigma)$.

It is easily checked that there are only $O(n^5)$ pairs of parallel rays for which at least one of the rays satisfies the first condition (and the other ray satisfies either condition). Hence, it suffices to bound the number of pairs of parallel rays both satisfying the second condition, i.e., each of them passes through three edges of Σ before intersecting the open region lying below Σ . Let ρ be such a ray, and let c be the third edge of Σ intersected by ρ . Then the ray emanating from the intersection point of c and ρ in the direction opposite to that of ρ passes three edges of Σ and is disjoint from the open region below Σ .

Suppose ρ_1 and ρ_2 are two such (reversed) rays, and suppose that ρ_1 passes through the edges a_1, b_1 , and c_1 of Σ in this order, and that ρ_2 passes through the edges a_2, b_2 , and c_2 in this order (note that the corresponding orthographic view is seen in the direction $(\theta + \pi, \pi - \varphi)$ opposite to the orientation (θ, φ) of these rays). Let us consider the fixed pair a_1, a_2 of edges of the terrain, and define the following collection $\mathcal{F}_{a_1 a_2}$ of trivariate partial functions. The three independent variables are (s_1, s_2, θ) , and each such triple represents the vertical projection of a pair of parallel rays, as follows: s_1 (resp. s_2) parametrizes a point on a_1 (resp. on a_2) that we denote by $p(s_1)$ (resp. $p(s_2)$). To be more precise s_1 is the distance between the left endpoint of a_1 (resp. a_2) and $p(s_1)$ (resp. $p(s_2)$). Finally, θ is the common horizontal orientation of the two rays, one of which emanates from the point $p(s_1)$ on a_1 and the other emanates from the point $p(s_2)$ on a_2 (the only variable degree of freedom available for these rays is their common azimuth φ).

For each edge e of Σ we define a function $F_e \in \mathcal{F}_{a_1 a_2}$ so that $F_e(s_1, s_2, \theta)$ is the smallest of the two azimuth angles φ_1, φ_2 , for which two rays, ρ_1, ρ_2 , exist at orientations (θ, φ_1) and (θ, φ_2) , respectively, so that ρ_1 (resp. ρ_2) emanates from $p(s_1)$ (resp. $p(s_2)$) and passes through e . If only one of these rays exists, $F_e(s_1, s_2, \theta)$ is equal to the azimuth of that ray, and if none of these rays exist, $F_e(s_1, s_2, \theta)$ is not defined.

Let F denote the lower envelope of the collection $\mathcal{F}_{a_1 a_2}$. For each triple (s_1, s_2, θ) the following is easily seen to hold: let $\varphi = F(s_1, s_2, \theta)$ (assuming that it is defined); then two parallel rays, with common orientation (θ, φ) , exist emanating respectively from $p(s_1) \in a_1$ and from $p(s_2) \in a_2$, so that both rays are disjoint from the region lying below Σ and at least one of them passes through another edge of Σ . (Some care must be exercised in handling the unbounded faces over which these rays pass, to ensure that they do not pierce through such a face. This can be enforced by adding extra functions to the collection $\mathcal{F}_{a_1 a_2}$; we leave it to the reader to work out the details.) Moreover, if $(s_1, s_2, \theta, \varphi)$ is a breakpoint of F , where four surfaces of $\mathcal{F}_{a_1 a_2}$ meet, then the two corresponding rays collectively pass through four edges of Σ (in addition to a_1 and a_2). Conversely, any pair of rays with these properties appear as a breakpoint of F , so, in particular, each critical orthographic view which is a vertex of $M(\Sigma)$ must give rise to such a breakpoint.

Hence, by the results of [11] (using an appropriate parametrization which makes the functions of $\mathcal{F}_{a_1 a_2}$ algebraic of constant maximum degree), the number of such critical pairs of rays is $O(n^{3+\epsilon})$. Multiplying this bound by the $O(n^2)$ pairs

of edges a_1, a_2 , we obtain that the overall number of such critical pairs of rays, and thus also the number of topologically different orthographic views of Σ , is $O(n^{5+\epsilon})$, for any $\epsilon > 0$. \square

Remark. (i) de Berg *et al.* [1] have given a lower bound of $\Omega(n^5\alpha(n))$ for the number of topologically different orthographic views of a polyhedral terrain with n edges. Hence, the above upper bound is close to optimal.

(ii) Theorem 1 can be extended to more general terrains. For example, Theorem 1 holds even if each face of Σ is a portion of a sphere (or of any constant-degree algebraic surface). In this case it suffices to bound the number of pairs of parallel rays, each of which is tangent to three faces of Σ and is disjoint from the open region lying below Σ . Following a similar approach it can be shown that the number of such pairs of parallel rays, and thus the number of topologically different orthographic views of such a terrain, is $O(n^{5+\epsilon})$, for any $\epsilon > 0$.

3. The Number of Perspective Views

Theorem 2. *The number of topologically different perspective views of a polyhedral terrain with n edges is $O(n^{8+\epsilon})$, for any $\epsilon > 0$.*

Proof. The proof is very similar to that of Theorem 1. Specifically, the analysis of [1] and the proof of Theorem 1 imply that it suffices to bound the number of points (x, y, z) for which three segments emanating from (x, y, z) exist, each of which passes through three edges of Σ and is disjoint from the open region lying below Σ . (Such a point (x, y, z) is an intersection between three surfaces of critical viewpoints in 3-space, where each surface is the locus of viewpoints at which some topological changes in the perspective view take place.) Suppose that ρ_1, ρ_2 , and ρ_3 are these segments, and suppose that ρ_i passes through the edges a_i, b_i , and c_i of Σ in this order when directed *toward* (x, y, z) , for $i = 1, 2, 3$. Let us consider the fixed triple a_1, a_2, a_3 of edges of Σ , and define the following collection $\mathcal{F}_{a_1a_2a_3}$ of 5-variate partial functions. The five independent variables are (s_1, s_2, s_3, x, y) , and each such 5-tuple represents the vertical projection of a triple of concurrent segments, as follows: s_i parametrizes a point on a_i , for $i = 1, 2, 3$, that we denote by $p(s_i)$ (see Section 2 for a more precise definition of s_1, s_2), and (x, y) is the vertical projection of the common endpoint of the three segments which connect this point to each of the three points $p(s_i)$, for $i = 1, 2, 3$.

For each edge e of Σ we define a function $F_e \in \mathcal{F}_{a_1a_2a_3}$ so that $F_e(s_1, s_2, s_3, x, y)$ is the largest z coordinate of a point $q = (x, y, z)$ with the property that at least one of the three segments connecting q to $p(s_1), p(s_2), p(s_3)$ passes through e (if none of these segments exists, $F_e(s_1, s_2, s_3, x, y)$ is undefined).

Let F denote the upper envelope of the collection $\mathcal{F}_{a_1a_2a_3}$. For each 5-tuple (s_1, s_2, s_3, x, y) the following is easily seen to hold: let $z = F(s_1, s_2, s_3, x, y)$ (assuming that it is defined); then the three segments connecting (x, y, z) to $p(s_1), p(s_2), p(s_3)$ are disjoint from the open region lying below Σ and at least one of them

touches another edge of Σ (again, we need to modify the collection of functions so as to handle properly the unbounded faces of Σ ; we omit the straightforward details). Moreover, if (s_1, s_2, s_3, x, y, z) is a breakpoint of F , where six surfaces of $\mathcal{F}_{a_1, a_2, a_3}$ meet, then the three corresponding segments collectively pass through six edges of Σ (in addition to a_1, a_2 , and a_3). Conversely, any triple of segments with these properties corresponds to a breakpoint of F .

Hence, by the results of [11], the number of such critical triples of segments is $O(n^{5+\varepsilon})$, where the constant of proportionality depends on ε . Multiplying this bound by the $O(n^3)$ triples of edges a_1, a_2, a_3 , we obtain that the overall number of such critical triples of segments, and thus also the number of topologically different perspective views of Σ , is $O(n^{8+\varepsilon})$, for any $\varepsilon > 0$. \square

Remark. (i) de Berg *et al.* [1] have given a lower bound of $\Omega(n^8 \alpha(n))$ for the number of topologically different perspective views of a terrain with n edges. Thus our bound is close to optimal in the worst case.

(ii) Unfortunately, both for orthographic and for perspective views, our technique does not yield an efficient algorithm for computing the aspect graph of a polyhedral terrain, whose running time is close to the maximum size of that graph, because no such efficient algorithm is known for computing the lower envelope of surface patches in dimensions greater than three.

References

1. M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld, Sparse arrangements and the number of views of polyhedral scenes, Manuscript, 1991.
2. K. Bowyer and C. Dyer, Aspect graphs: An introduction and survey of recent results, *Internat. J. Imaging Systems Technol.* **2** (1990), 315–328.
3. I. Chakravarty and H. Freeman, Characteristic views as a basis for three-dimensional object recognition, *Proc. SPIE: Robot Vision* **336** (1982), 37–45.
4. R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *J. Symbolic Comput.* **7** (1990), 11–30.
5. C. Crawford, Aspect graphs and robot vision, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, 1985, pp. 382–384.
6. Z. Gigus, J. Canny, and R. Seidel, Efficiently computing and representing aspect graphs of polyhedral objects, *IEEE Trans. Pattern Anal. Mach. Intell.* **13** (1991), 542–551.
7. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions with applications to visibility in terrains, *Proc. 9th ACM Symp. on Computational Geometry*, 1993, pp. 11–18.
8. J. Koenderink and A. van Doorn, The singularities of visual mapping, *Biol. Cybernat.* **24** (1976), 51–59.
9. W. Plantinga and C. Dyer, An algorithm for constructing the aspect graph, *Proc. 27th IEEE Symp. on Foundations of Computer Science*, 1986, pp. 123–131.
10. W. Plantinga and C. Dyer, Visibility, occlusion, and the aspect graph, *Internat. J. Comput. Vision* **5** (1990), 137–160.
11. M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Proc. 34th IEEE Symp. on Foundations of Computer Science*, 1993, pp. 498–507.
12. J. Snoeyink, The number of views of axis-parallel objects, *Algorithms Rev.* **2** (1991), 27–32.

Received March 30, 1993, and in revised form June 28, 1993.