Sparse Arrangements and the Number of Views of Polyhedral Scenes*

Mark de Berg† Dan Halperin‡ Mark H. Overmars† Marc van Kreveld†

Abstract

In this paper we study several instances of the problem of determining the maximum number of topologically distinct two-dimensional images that three-dimensional scenes can induce. To bound this number, we investigate arrangements of curves and of surfaces that have a certain sparseness property. Given a collection of n algebraic surface patches of constant maximum degree in 3-space with the property that any vertical line stabs at most k of them, we show that the maximum combinatorial complexity of the entire arrangement that they induce is $\Theta(n^2 k)$. We extend this result to collections of hypersurfaces in 4-space and to collections of $(d-1)$-simplices in d-space, for any fixed d. We show that this type of arrangements (*sparse* arrangements) is relevant to the study of the maximum number of topologically different views of a polyhedral terrain. Given a polyhedral terrain with n edges and vertices, we derive an upper bound $O(n^5 \cdot 2^{c\sqrt{\log n}})$ on the maximum number of views of the terrain from infinity, where c is some positive constant. Moreover, we show that this bound is almost tight in the worst case, by introducing a lower bound construction inducing $\Omega(n^5\alpha(n))$ distinct views. We also analyze the case of perspective views, point to the potential role of sparse arrangements in obtaining a sharp bound for this case, and present a lower bound construction inducing $\Omega(n^8\alpha(n))$ distinct views.

For the number of views of a collection of k convex polyhedra with a total of n faces, we show a bound of $O(n^4k^2)$ for views from infinity and $O(n^6k^3)$ for perspective views. We also present lower bound constructions for such scenes, with $\Omega(n^4 + n^2k^4)$ distinct views from infinity and $\Omega(n^6 + n^3k^6)$ views when the viewpoint can be anywhere in 3-space.

1 Introduction

In this paper we study several instances of the so-called *aspect graph* problem, which has recently attracted much attention, especially in computer vision. Aspect graphs are often

*This research was supported by the ESPRIT Basic Research Action No. 3075 (project ALCOM). The first and third authors were also supported by the Dutch Organization for Scientific Research (N.W.O.).

†Department of Computer Science, Utrecht University, P.O.Box 80.089, 3508 TB Utrecht, the Netherlands.

‡Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305. Work on this paper was carried out while this author was visiting Utrecht University.

studied in the context of three-dimensional scene analysis and object recognition. The complexity of an aspect graph is determined by the number of topologically different views of a scene. To bound this number, we investigate arrangements of curves and of surfaces that have a certain sparseness property.

1.1 Background

At a high level, the aspect-graph[1] problem can be formulated as follows: Given a three-dimensional scene consisting of one or more three-dimensional objects, how many qualitatively different (see below for a precise definition) two-dimensional images can the scene induce and how efficiently can one compute and represent a partitioning of the *viewing space*, i.e., the space of possible placements of the viewpoint, into maximal connected portions having the same view each. A concrete instance of the problem is specified by determining (i) the type of objects in the scene; (ii) the viewing space; and (iii) what makes a pair of images of the scene (qualitatively) different. Koenderink and van Doorn are credited for introducing this concept [15],[16]. Since then, aspect graphs have attracted a lot of interest, mainly in the computer vision community, e.g., [5], [6], [14], [17], [24], [25]. We also mention the works of Plantinga and Dyer [19] and of Gigus et al. [11] that have a computational geometry flavor.

Here we give the basic terminology needed in the sequel. For a broader introduction and a survey of recent research on aspect graphs see, e.g., [4], from which we borrow most of the subsequent terminology. In this paper, we restrict ourselves to polyhedral scenes, where every face of an object is flat and any induced image of an object is a straight line drawing. We also assume that the objects in the scene are *opaque*. We will consider two types of viewing spaces. One type is the collection of views from infinity which can be modeled by the *sphere of directions*. Conceptually, we place a large sphere centered at the origin around our scene, and each point on the surface of the sphere represents the direction of view from that point towards the origin. For every direction, the view of the scene is the result of an orthographic (parallel) projection of the visible portions of the objects in the scene in a direction opposite to the viewing direction onto a plane which is far from the scene and is orthogonal to the viewing direction. The second and more general viewing space is where we allow the viewing point to be anywhere in the 3D space of the scene, and a view from a point p is the perspective projection of the scene as seen from p. The perspective view of a scene from a point p can be illustrated by considering an infinitesimally small sphere centered at p, onto which the scene is projected.

A fixed orthographic view of the scene can be regarded as a straight-edge planar subdivision consisting of faces, edges and vertices. As the viewing direction changes continuously, the view (i.e., the induced planar subdivision) changes continuously until we reach a critical direction at which the subdivision undergoes a topological or combinatorial change: vertices, edges or faces newly appear or disappear. (In compliance with other work in the area, we will use *topological* to refer to these type of changes.) The loci of critical directions

[1]The term *aspect graph* is synonymous with *characteristic views, viewing data* and other similar terms.

lie on curves on the sphere of directions that we will refer to as *critical curves*. Similarly, a fixed perspective view of the scene can be regarded as a subdivision of a sphere consisting of faces, circular-arc-edges and vertices. We consider two views to be the same if the topological structure of their respective subdivisions is the same [20].

We wish to partition the viewing space into maximal connected regions such that inside one region all the views are the same. For views from infinity, we aim to partition the sphere of directions into maximal faces having the same view topologically (these will be referred to as *general viewpoints* [4]), separated by critical curves, which represent *accidental viewpoints*. (See Section 3 for a detailed discussion of these curves and the partitioning they induce.) For perspective views we aim to partition the entire space of the scene into maximal connected 3D regions having the same view and separated by critical surfaces.

The term *aspect graph* originates from a certain representation of the viewing space as a discrete graph where each node of the graph represents a maximal connected component of the space having the same aspect (or view). Plantinga and Dyer [19] have shown that the maximum number of views of a convex polyhedron with n vertices is $\Theta(n^2)$ for views from infinity and $\Theta(n^3)$ for perspective views. Later, it has been shown that for a general polyhedron, or more generally, for a collection of n non-intersecting triangles in space, the maximum number of views can be as high as $\Theta(n^6)$ for orthographic views and $\Theta(n^9)$ for perspective views [20]. Snoeyink [23] has shown that even if we restrict the objects to be axis-parallel polyhedra, the bound for orthographic views remains $\Theta(n^6)$.

1.2 Summary of Results

In this paper we study the following instances of the aspect-graph problem, where better bounds can be shown: (i) The case where the scene consists of a *polyhedral terrain* with a total of n edges; and (ii) the case where the scene consists of k convex polyhedra with a total of n edges. A polyhedral terrain is the graph of a piecewise-linear (polyhedral) continuous function $z = F(x, y)$ defined over the entire xy-plane. Cole and Sharir [9] have studied a variety of visibility problems for polyhedral terrains, and showed that the maximum number of distinct views when the viewpoint moves along a fixed vertical line is considerably smaller than the number of distinct views when the viewpoint moves along a line in any other direction.

To bound the overall number of views of a polyhedral terrain, we need additional machinery and we consider a special type of *arrangements*. An arrangement of surfaces in d-space is the partitioning of d-space induced by a collection of surfaces. Arrangements play a central role in computational geometry, and the analysis of many geometric algorithms relies on the *complexity* of an arrangement or of portions of an arrangement (see, e.g., [10], [12]). The complexity of an arrangement of surfaces in 3-space, for example, is the overall number of faces of dimensions $0, 1, 2$ and 3 in the partitioning of space induced by these surfaces. We obtain the following result which we believe to be of independent interest (Proposition 2.3):

Given a collection of n surface patches in three-dimensional space all algebraic of constant maximum degree and bounded by a small number (bounded by a constant) of algebraic arcs of constant maximum degree, with the additional property that every vertical line stabs at most k of the surface patches, $k > 1$, the arrangement induced by these surface patches has complexity $\Theta(n^2 k)$.

This generalizes and improves a result by Sharir [22], who gives an $O(n^2 k\alpha(n/k))$ bound for the case of triangles[2]. We generalize the result even further to collections of hypersurfaces in 4-space and collections of $(d-1)$-simplices in d-space with a low 'vertical stabbing number'.

Using an analogous result in the plane we show that the maximum number of views of a terrain with n vertices is $O(n^5 \cdot 2^{c\sqrt{\log n}})$ for views from infinity, for some positive constant c. Furthermore, we show that this bound is almost tight in the worst case by presenting a polyhedral terrain with n edges that induces $\Omega(n^5\alpha(n))$ distinct orthographic views. We then turn to analyze the case of perspective views and point to a potential use of the result for sparse arrangement of surfaces in 3-space, to obtain a sharp bound for this case. More precisely, we show that the bound on the number of perspective views of a terrain with n vertices is $O((\nu(n))^2 n\lambda_4(n))$ where $\nu(n)$ is the maximum complexity of a certain family of segments defined relative to a terrain (see Section 3.2 for more details), and where $\lambda_4(n)$ is a near-linear function related to Davenport-Schinzel sequences, $\lambda_4(n) = \Theta(n2^{\alpha(n)})$ [2]. We then present a lower bound $\Omega(n^8\alpha(n))$ for this quantity. We also investigate arrangements of k convex polyhedra having a total of n faces—where *any* line stabs at most $2k$ faces—and we obtain an improved and tight bound $\Theta(nk^2)$ on the maximum complexity of such an arrangement.

Finally, we study another instance of the aspect-graph problem where the scene consists of k opaque convex polyhedra having a total of n faces. In this case we show that the number of curves (or alternatively surfaces) determining the partitioning of the view space is only $O(n^2 k)$ (instead of $\Theta(n^3)$ in the general case) and we obtain a bound $O(n^4 k^2)$ on the maximum number of views from infinity and $O(n^6 k^3)$ for perspective views. For this type of scenes, we present constructions that induce $\Omega(n^4 + n^2 k^4)$ distinct views from infinity and $\Omega(n^6 + n^3 k^6)$ views when the viewpoint can be anywhere in space.

Two papers related to the study in this paper have recently been published. One paper, by Halperin and Sharir [13] uses part of the analysis given in Subsection 3.1 below in combination with new results to demonstrate the applicability of the main results of [13]. In this sense, Theorem 3.2 below is a joint result of both studies. In another paper, Agarwal and Sharir [1] devise an alternative technique to bound the number of views of a polyhedral terrain. For orthographic views, their technique produces an upper bound that is inferior to the bound given below. For perspective views they obtain a bound $O(n^{8+\varepsilon})$ which is almost tight in the worst case (as our lower bound construction in Subsection 3.2 shows), and is the best known (see also Remark 3.5 below).

The paper is organized as follows: In Section 2 we derive a collection of combinatorial results concerning sparse arrangements in two-, three- and higher dimensions. We then

[2] Here and throughout the paper, $\alpha(n)$ is the extremely slowly growing functional inverse of Ackermann's function.

apply some of these results, in Section 3, to obtain bounds on the maximum number of views of polyhedral terrains. In Section 4 we consider arrangements of convex polyhedra. In Section 5 we bound the maximum number of views of collections of convex polyhedra. Some concluding remarks and open problems are presented in Section 6.

2 Arrangements of Surfaces with Low Vertical Stabbing Number

This section deals with arrangements of surfaces where any vertical line intersects only a bounded size subset of the surfaces. In Subsection 2.1 we obtain several combinatorial results for the two- and three-dimensional cases that we will be using in the next section. In Subsection 2.2, we extend these results to arrangements of hypersurfaces in 4-space and to arrangements of $(d-1)$-simplices in d-space, for any fixed d.

2.1 Combinatorial Analysis

We start with the easier case of arrangements of curves in the plane and then proceed to handle arrangements of surfaces in 3-space.

Consider an arrangement of n simple curves in the plane, where a pair of curves intersect at most s times for some constant s. The maximum complexity of the entire arrangement in such a case is clearly $\Theta(n^2)$. We are interested in arrangements of curves that have the additional property that every vertical line intersects the curves in a total of at most k points.[3] The following result has been previously obtained by several authors (we are aware of a simple and tight bound by Pach, and an almost tight bound by Sharir—both can be found in [22]). We present another simple proof that gives a tight bound. Later we will use a generalization of it for the three-dimensional case.

Lemma 2.1 *Given a collection of n Jordan arcs in the plane, where every pair intersect at most a constant number of times and any vertical line stabs the arcs in a total of at most k points, then the maximum complexity, $B(n,k)$, of the partitioning of the plane induced by these curves is $\Theta(nk)$.*

Proof. Partition the plane into n/k vertical slabs such that each slab contains at most $2k$ endpoints of the curves. Inside each slab we have at most $2k$ curves: We consider the intersection of a curve with the vertical boundary of the slab as an endpoint; we thus have at most $4k$ potential endpoints at our disposal—$2k$ inside the slab and $2k$ on its boundaries, therefore we can "pay" for at most $2k$ curves. Hence, there are at most $O(k^2)$ intersection points inside each slab. The total number of intersection points is therefore $n/k \cdot O(k^2) = O(nk)$. A bound on the maximum number of intersection points obviously serves as an (asymptotic) upper bound on the complexity of the arrangement.

[3] Since for the two-dimensional case we do *not* require the curves to be algebraic of constant maximum degree, the actual number of intersection points with the curves counts.

The lower bound follows from the lower bound in Proposition 2.4 (presented below) with $d = 2$. $\square$

Next, we consider arrangements of algebraic surface patches (2-manifolds with boundary) in three-dimensional space. We assume the surface patches that we deal with to be algebraic of maximum degree b, where b is a small constant. Also, we assume that the boundary of each surface patch consists of a small constant number of algebraic curves, all of maximum degree, say b too. There are a few ways to extend the two-dimensional problem to the three-dimensional case. A straightforward extension is the following:

Lemma 2.2 *Given a collection of n algebraic surface patches of constant maximum degree in 3-space such that any plane parallel to the yz plane intersects only k of them, then the maximum complexity, $B'(n,k)$, of the entire arrangement induced by these surface patches is $\Theta(nk^2)$.*

But for our purposes (as will be discussed in the next section) we need a different extension whose proof requires the use of a more powerful divide-and-conquer technique.

Proposition 2.3 *Given a collection of n algebraic surface patches of constant maximum degree in three-dimensional space such that any vertical line stabs at most k of them, $k > 1$, then the maximum complexity, $D(n,k)$, of the arrangement induced by these surface patches is $\Theta(n^2k)$.*

Proof. First we decompose each surface patch into a constant number of surface patches, with the property that any vertical line intersects any patch in at most one point. We denote the resulting collection of surface patches by S. Then we project the boundaries of the patches onto the xy-plane. This gives a set C of $O(n)$ algebraic curves of constant maximum degree (where the constant may be higher than the constant bounding the degree of the original surfaces).

Next, we use random sampling (see [8]) to control the divide-and-conquer process. We choose a random subset of curves $R \subset C$ of size $r \leq n$, where each r-element is chosen with equal probability, and consider the arrangement $\mathcal{A}(R)$, which admits a vertical decomposition into $m = O(r^2)$ faces $f_1, f_2, \ldots, f_m$. Let n_i be the number of curves in C crossing the face f_i. From the analysis of Clarkson and Shor [8], it follows that for any fixed integer $\nu \geq 0$, the expected value of $\sum_{i=1}^{m} n_i^\nu$ is $O(r^2(n/r)^\nu)$. We choose a sample R of size $r \leq n$ for which $\sum_{i=1}^{m} n_i^3 = O(n^3/r)$. Consider one face f_j in the decomposition of the arrangement $\mathcal{A}(R)$ and let S_1 be the subset of surfaces of S whose projection onto the xy-plane fully contains the face f_j. Let S_2 be the subset of surfaces of S for which the projection of their 1D boundary crosses f_j.

By the assumption of low vertical stabbing number, we know that $|S_1| < O(k)$. By definition $|S_2| = n_j$. Therefore, the complexity of the arrangement above the face f_j is $O((k+n_j)^3)$. Hence

$$D(n,k) = O(\sum_{i=1}^{m}(k+n_i)^3) .$$

Choosing $r = \lceil \frac{n}{k} \rceil$ leads to the desired bound

$$D(n, k) = O(n^2 k).$$

That the bound is tight follows from the lower bound in Proposition 2.4 with $d = 3$. □

Obviously, if $k = 1$, then the surfaces are pairwise disjoint and therefore the complexity of the entire arrangement is $\Theta(n)$. See also Remark 2.5 below.

2.2 Extension to Higher Dimensions

The proof of the previous results relies on "good" partitioning schemes in $(d-1)$-dimensional space. Such partitionings are available for arrangements of simplices in any fixed dimension and for arrangements of low-degree algebraic surfaces in 3-space.

Proposition 2.4 *Given a collection of n $(d-1)$-simplices in E^d for a fixed d, such that any vertical line (i.e., a line parallel to the X_d axis) stabs at most k of them, then the arrangement induced by these simplices has maximum complexity $O(n^{d-1}k)$. Furthermore, this bound is tight for $k > d - 2$.*

Proof. Project the simplices onto the hyperplane $X_d = 0$ and construct a $(1/r)$-cutting[4] of size $O(r^{d-1})$ for the hyperplanes supporting the boundaries of the projections of the simplices (see [18]). Taking $r = \lceil \frac{n}{k} \rceil$ we can bound the complexity $D_d(n, k)$ of the entire arrangement as follows:

$$D_d(n, k) = O(r^{d-1}(k + n/r)^d) = O(n^{d-1}k).$$

For the lower bound, construct a "grid" made of n $(d-2)$-simplices on the hyperplane $X_d = 0$ that has complexity $\Omega(n^{d-1})$. Extend each $(d-2)$-simplex in the X_d direction into a "long" $(d-1)$-simplex. Finally, cut the resulting $(d-1)$-simplices by additional $k - d + 1$ $(d-1)$-simplices, all parallel to the hyperplane $X_d = 0$. See Figure 1 for an illustration of the construction in 3-space. (In the figure we use, for convenience, rectangles rather than triangles, but these can easily be replaced by triangles without affecting the lower bound.) □

Remark 2.5 The reason why the above lower bound holds only for $k > d - 2$ is that the grid of the construction has edges each of which is the intersection of $d-1$ $(d-1)$-simplices. Consequently, the grid itself requires that the vertical stabbing number be at least $d - 1$.

Proposition 2.6 *Given a collection of n algebraic hypersurfaces of constant maximum degree in four-dimensional space such that any vertical line (i.e., a line parallel to the X_4 axis) stabs at most k of them, the arrangement induced by these surfaces has complexity $O(n^3 k\, \beta(\frac{n}{k}))$, where $\beta(\cdot)$ is an extremely slowly growing function.*[5] *Moreover, for $k > 2$, the*

[4] Given a set S of n $(d-1)$-simplices in E^d, a $(1/r)$-*cutting* for S is a collection Ξ of (possibly unbounded) closed d-simplices which together cover E^d and such that the interior of each simplex in Ξ is intersected by at most $\frac{n}{r}$ $(d-1)$-simplices of S. For more details, see, e.g., [18].

[5] The function $\beta(n)$ is defined in [7]: $\beta(n) = 2^{\alpha(n)^c}$, where c is a constant depending on the degree of the surfaces that are the projection of the original hypersurfaces onto the hyperplane $X_4 = 0$.

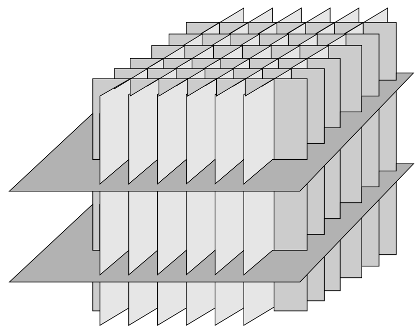

Figure 1: A three-dimensional arrangement with $\Theta(n^2k)$ complexity

complexity can be as large as $\Omega(n^3k)$.

Proof. The proof is similar to the proof of Proposition 2.3, and it uses random sampling and the stratification scheme of Chazelle et al. [7]. The lower bound follows from the lower bound of Proposition 2.4 for the case $d = 4$. □

3 The Number of Views of Polyhedral Terrains

In this section, we show how sparse arrangements are relevant to the analysis of the number of views of polyhedral terrains. We obtain an upper bound on the maximum number of views of polyhedral terrains when viewed from infinity, and we show that this bound is almost tight in the worst case, by introducing a lower bound construction that almost achieves the upper bound. We also analyze the case of perspective views, and point to a potential role of sparse arrangements in obtaining a good bound for this case. Finally, we present a lower bound construction for perspective views.

A polyhedral terrain is the graph of a piecewise-linear (polyhedral) continuous function $z = F(x, y)$ defined over the entire xy-plane. We assume that the graph has n edges. Since the projection of the terrain onto the xy-plane is a planar map, the number of vertices and faces of the polyhedral terrain is $O(n)$. Cole and Sharir [9] study a variety of visibility problems for polyhedral terrains. In particular, they consider the number of views of a terrain when the viewpoint is restricted to move along a given vertical line. We will be using their result, which we now state (in a slightly modified manner):

Theorem 3.1 (Cole and Sharir [9]) *The maximum number of different views of a polyhedral terrain with* n *edges when the viewpoint moves along a given vertical line is* $O(n\lambda_4(n))$, *that is* $O(n^2 2^{\alpha(n)})$.

We extend this result to a larger view space—the space of views from infinity. In the following subsection we handle the views from infinity and in Subsection 3.2 we discuss perspective views.

3.1 Views from Infinity

To bound the number of views from infinity we partition the *sphere of directions* into maximal connected components such that the view from any two points inside one component is topologically the same. Our goal is to obtain a bound on the maximum number of these components. The partitioning is induced by curves of three types (see [11] for a detailed study of these curves):

I. A curve defined by the plane through a face of the terrain—this curve is a great circle on the sphere of directions which is the intersection with the plane through the center of the sphere of directions that is parallel to the face.

II. A curve defined by a vertex-edge pair of the terrain. It is also a great circle on the sphere of directions resulting from intersecting the sphere of directions with a plane through the center that is parallel to the plane that passes through the vertex and the edge.

III. A curve describing the union of directions of lines that pass through the same three edges of the terrain. More specifically, for a fixed triple of edges of the terrain, consider the collection of lines that pass through these three edges. Translate all lines in this collection to contain the origin, and let β be the surface that is the union of the translated lines. The curve of type III is the intersection of β with the sphere of directions.

By definition, there are $O(n)$ curves of the first type, $O(n^2)$ curves of the second type and $O(n^3)$ curves of the third type. These are all algebraic curves of low degree [11]. An immediate, naive bound on the number of different views is $O(n^6)$ which is the maximum number of faces in a partitioning of a plane (or a sphere) by $O(n^3)$ curves, each pair of which does not intersect more than some constant number of times. But for a polyhedral terrain, we can obtain an improved bound.

For our purposes we need a more refined analysis of the curves of type III. In order to be able to use Theorem 3.1 in our case, we need to distinguish between visible and invisible portions of curves of type III. We substitute each curve $\gamma = \gamma(e_1, e_2, e_3)$ of type III defined by the edges e_1, e_2, e_3 of the terrain, 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 the terrain or else penetrates the terrain only at points that lie further away from its contacts with the edges e_1, e_2, e_3. In other words, for a point p on a curve γ to be visible, we require the existence of a ray whose direction is opposite to the viewing direction represented by p, that touches e_1, e_2, and e_3 but otherwise lies fully above the terrain. 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 touches the terrain at four edges before penetrating the terrain.

It has been recently shown [13] that for a terrain with n edges, the overall number of such rays is at most $O(n^3 \cdot 2^{c\sqrt{\log n}})$, for some absolute positive constant c. In summary, the maximum number of maximal visible portions of critical curves of the third type is $O(n^3 \cdot 2^{c\sqrt{\log n}})$. (Indeed, part of the analysis in this subsection is used by Halperin and Sharir [13] to demonstrate the applicability of the main results in [13].) We are now ready to prove the following:

Theorem 3.2 *The maximum number of topologically distinct views of a polyhedral terrain with a total of n edges, when viewed from infinity, is*

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

for a positive constant c' (c' slightly larger than c).

Proof. We assume, without loss of generality, that the terrain has a minimum z-value $z = z_0$. We fix the center of the sphere of directions to lie on the plane $z = z_0$ and so we are only interested in the upper hemisphere. It is not difficult to see that in terms of views from infinity, Theorem 3.1 can be rephrased to give the bound $O(n\lambda_4(n))$ on the maximum number of views when letting the viewpoint move along a fixed meridian on the sphere of directions. (Here, by a *meridian*, we refer to the portion of a great circle through the poles that lies between the equator and the "north" pole.) This implies that as we move the viewpoint along the meridian, although there are $O(n^3 \cdot 2^{c\sqrt{\log n}})$ curves on the sphere, it does not cross more than $O(n\lambda_4(n))$ curves on its way. The adaptation of Lemma 2.1 from the planar case to our case is immediate, and therefore we may now employ Lemma 2.1 where the total number of curves in this case is $O(n^3 \cdot 2^{c\sqrt{\log n}})$ and the "vertical" stabbing number is $O(n\lambda_4(n))$, and the bound follows. □

We now turn to discuss a lower bound for the number of orthographic views of a terrain. Our approach to designing lower bound constructions for the number of views consists of building a separate construction for every degree of freedom of the viewpoint such that when fixing one degree of freedom, all the views for the other degree(s) of freedom are attainable. Thus, the number of views of the whole construction is the product of the number of views for each degree of freedom. (This approach will be further exemplified in Section 5.)

Theorem 3.3 *There exists a polyhedral terrain with n edges for which the number of distinct orthographic views is $\Omega(n^5\alpha(n))$.*

Proof. For views from infinity, we may regard our viewpoint as moving on the sphere of directions, that is, it has two degrees of freedom. Therefore, our construction consists of two sub-constructions. Since our construction will use viewpoints belonging to only a certain portion of the sphere of directions, we may think of our viewpoint being located at a certain area far away from the scene, and therefore we can employ terms like *nearer to* or *further away from* the viewpoint.

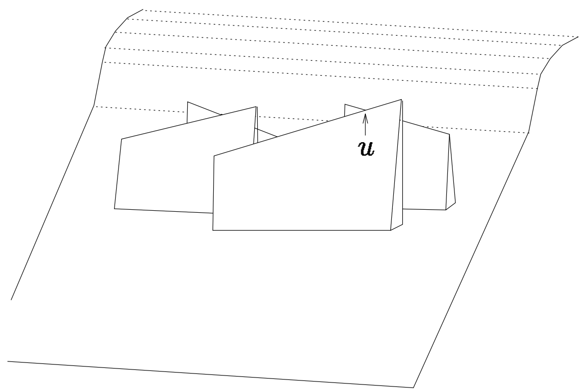

Figure 2: Construction A induces $\Omega(n^2\alpha(n))$ different views when the viewpoint moves along a meridian

Construction A for views that change when the viewpoint moves along a meridian on the sphere of directions, we adapt from [9] (see Figure 2): We take n segments whose upper envelope complexity is $\Omega(n\alpha(n))$ (as in [26]). These segments are put on parallel vertical planes and a thin wedge is drawn downwards from each of them. Further away from the viewpoint we construct a "hill" consisting of parallel horizontal slabs (i.e., the edges of the hill are parallel to the xy-plane). The upper envelope of the thin wedges is constructed such that when moving the viewpoint along a meridian, each view where a vertex of the upper envelope coincides with an edge of the hill happens at a distinct point along the meridian. This can be achieved for example, by taking all the breakpoints of the upper envelope to lie very close to one horizontal line, so that when moving up along a meridian the last time a vertex of the upper envelope coincides with an edge e of the hill is before the first time a vertex of the upper envelope coincides with the edge lying immediately below e on the hill. Thus, while moving along a meridian, we get $\Omega(n^2\alpha(n))$ different views.

Construction B is for views that change when the viewpoint moves along a *parallel of latitude* (the intersection of a plane parallel to $z = 0$ with the sphere of directions) of the sphere. This construction is an adaptation of a part of a construction by Canny (reported in [20]) for the lower bound for the number of views of arbitrary polyhedra. Far from the viewpoint we construct a hill similar to the hill in the construction above. See Figure 3. In front of the hill, nearer to our viewpoint, we construct a collection of n pyramids in a row which we denote by S. So far we have created a slanted grid—for a fixed view, a vertex of the grid is created by the intersection (of the projection) of a visible edge of a pyramid and an edge of the hill. Finally, farther from the grid and nearer to the viewpoint we construct a collection of n almost flat prisms which we denote by S'. When viewed from our viewpoint, the prisms of S' resemble wide rectangles, and every edge of a prism that extends from the horizontal plane upwards is very steep, almost vertical. The distance between the adjacent

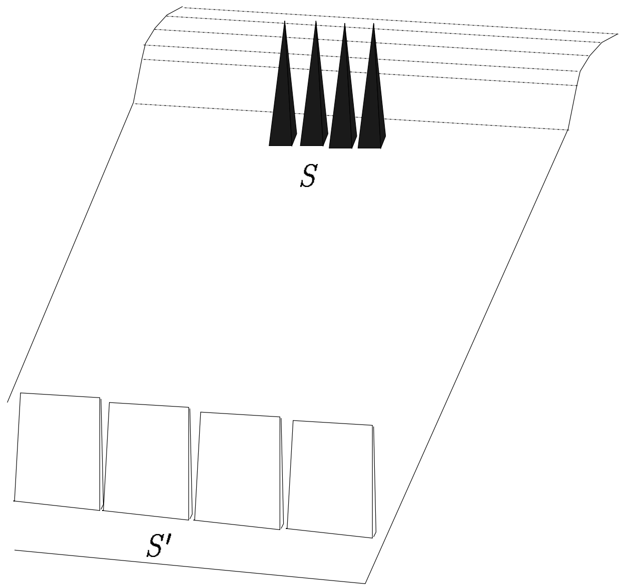

Figure 3: Construction B induces $\Omega(n^3)$ different views when the viewpoint moves along a parallel of latitude

quasi-vertical edges of two neighboring prisms is chosen to be very small. The distances are chosen such that when we move the viewpoint on a parallel of latitude, we see all the intersection points of the grid through one interval between a pair of prisms of S', before we see any other intersection point of the grid through another interval between another pair of prisms of S'. See Figure 4. Consider one such interval between a pair of prisms of S'. The edges that define this viewing "crack" are almost vertical, whereas the edges of the pyramids have a smaller slope. As we move the viewpoint on a parallel of latitude, each vertex of the grid will coincide with the, say, left edge of the interval at a different viewpoint, inducing $\Omega(n^2)$ different views for one interval. Since there are $\Omega(n)$ distinct intervals between adjacent prisms, we get $\Omega(n^3)$ distinct views when moving on a parallel of latitude.

We next interpret the two constructions in terms of critical curves on the sphere of directions. Construction A induces $\Omega(n^2\alpha(n))$ views when the viewpoint moves along a meridian. More precisely, it is placed such that these changes will be similar for a family of meridians, namely, there is a region of the sphere of directions where construction A induces a set of $\Omega(n^2\alpha(n))$ critical curves which are roughly parallel to the equator (or to any parallel of latitude). Similarly, construction B induces a set of $\Omega(n^3)$ critical curves that are roughly parallel to a meridian. The two constructions are juxtaposed such that these two sets of curves create a grid on the sphere of directions. Inside each two-dimensional face of the grid we get a distinct view. Figure 5 shows the final construction.

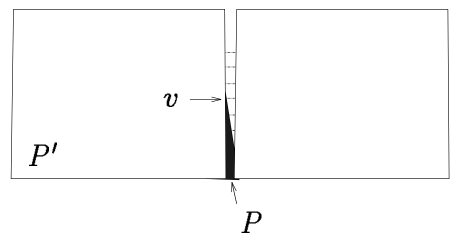

Figure 4: A view where three edges are coincident at the point v: an edge of the hill, an edge of a pyramid P in S and an almost vertical edge of a prism P' in S'

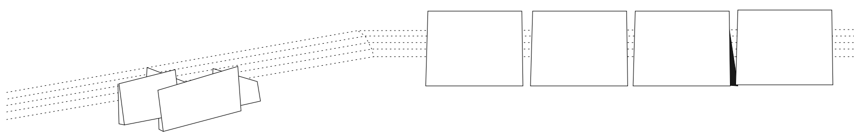

Figure 5: The overall construction inducing $\Omega(n^5\alpha(n))$ distinct orthographic views

How do we get the effect of roughly parallel critical curves in each of the constructions? Consider first construction A and specifically consider a vertex u in the upper envelope in the view presented in Figure 2. The vertex u is the meeting point of the projection of two edges e_1 and e_2 of the terrain. Because we are dealing with polyhedral terrains, these two edges cannot coincide, but we can make them arbitrarily close to one another. Thus, if we look at the critical curve of type III induced by e_1, e_2 and an edge of the hill behind them, we get a curve that locally resembles an arc of a parallel of latitude (as if the simultaneous view of e_1 and e_2 is actually a view of a vertex of the terrain).

Similarly, in construction B, we can make the pyramids in S arbitrarily close to the hill behind them, so that when we take an edge e_1 of the hill, an edge e_2 of a pyramid in S, and a quasi-vertical edge of a prism in S', the resulting critical curve of type III resembles a portion of a meridian (as if the simultaneous view of e_1 and e_2 is actually a view of a vertex of the terrain).

If we choose the proportions of the two constructions carefully—in particular we make the hill in construction A sufficiently long and the pyramids and prisms in construction B sufficiently high—then we get the desired grid effect on the sphere of directions. Therefore, in total the number of views of a polyhedral terrain when viewed from infinity can be $\Omega(n^5\alpha(n))$ in the worst case. $\square$

A small gap still remains between the upper bound shown in Theorem 3.2 and the lower bound shown in Theorem 3.3.

3.2 Perspective Views

For perspective views the viewpoint may be anywhere in 3-space. Surfaces similar to the surfaces that we have previously used to define curves on the viewing sphere now serve to partition 3-space into maximal connected (three-dimensional) cells where the perspective view does not change topologically. More precisely, we use the following surfaces: planes that contain faces of the terrain (type I), planes that pass through a vertex and an edge of the terrain (type II), and surfaces each of which is the union of lines that touch three fixed edges of the terrain simultaneously (type III). In this case we are unable to obtain a sharp bound as we have obtained for views from infinity. However, we point to the potential use of sparse arrangements here, and present a lower bound for perspective views. We also mention a recent result [1] where an upper bound for this case has been obtained using a different approach; see Remark 3.5 below.

By arguments similar to those we used above for views from infinity, there are $O(n^3)$ "critical" surfaces that subdivide the viewing space. Here also, we need a more refined analysis of the critical surfaces of type III (defined by triples of edges). For a fixed triple of edges, we replace the surface $\sigma = \sigma(e_1, e_2, e_3)$ by a collection of *visible* surface patches. A point in $\sigma(e_1, e_2, e_3)$ is said to be visible if there is a segment pq lying above the terrain which touches (e_1, e_2, e_3) and such that the viewpoint p does not lie on any of the three edges. A surface σ will be divided into patches, exactly where a corresponding line segment touches four edges of the terrain but does not penetrate the terrain on both sides. More precisely, this will occur when a line segment touches four edges of the terrain, one of its endpoints may lie on one of the edges, but the other endpoint should not lie on the terrain. This latter endpoint represents the viewpoint and hence it is not allowed to lie on the terrain.

The extra bounding curves that divide an original critical surface into patches, are all straight line segments, rays or lines. It is easily verified that we can subdivide the collection of visible portions of any such surface $\sigma(e_1, e_2, e_3)$ into a number of surface patches that is proportional to the number of extra bounding curves (by a standard two-dimensional vertical decomposition, for example). Hence, it is desirable to have a bound on the maximum number of extra bounding curves, which is bounded, in turn, by the maximum overall number of maximal segments that touch the terrain in four edges, lie above the terrain, and only one of the endpoints is allowed to lie on one of the four edges (for the reason we mention above, namely, one of the endpoints of each segment represents the viewpoint, and therefore we do not allow it to lie on a terrain edge). We denote the maximum complexity of this family of segments for any terrain with n edges, by $\nu(n)$.

Lemma 3.4 *The maximum number of combinatorially distinct perspective views of a polyhedral terrain with a total of n edges is $O((\nu(n))^2 n\lambda_4(n))$.*

Proof. We have a collection of $\nu(n)$ critical surface patches each with a small number (bounded by some constant) of bounding curves, which partition the viewing space into non-critical regions. Theorem 3.1 implies that when we let the viewpoint move along a fixed vertical line, it does not cross more than $O(n\lambda_4(n))$ of these surfaces. This is an upper bound on the vertical stabbing number of the arrangement of $O(\nu(n))$ surfaces. Plugging these quantities into Proposition 2.3 we get the asserted bound. □

The trivial upper bound on $\nu(n)$ is $O(n^4)$, as it is well-known that a segment cannot touch four edges of a terrain in more than two placements (assuming general position). However, for our purposes the goal is to show that $\nu(n)$ is roughly cubic, then the resulting bound on the number of perspective views will be roughly $O(n^8)$, the same order of magnitude as the lower bound that we present below. We remark that a lower bound of $\Omega(n^3)$ for $\nu(n)$ is easy to establish.

Remark 3.5 (1) Recently, Agarwal and Sharir [1] have shown an upper bound $O(n^{8+\varepsilon})$ for any $\varepsilon > 0$, on the maximum number of perspective views of a terrain with n edges, using lower envelopes of hypersurfaces in 6-dimensional space. They have also applied the same approach to obtain a bound $O(n^{5+\varepsilon})$ for the case of orthographic views. Thus our analysis of the orthographic case yields an improved bound.
(2) The result of [1] for perspective views almost settles the problem (see our lower bound construction below). However, if a sharp bound is obtained for $\nu(n)$, it might result in an improved bound for this case as well.

We conclude this section with a lower bound construction for the perspective case.

Lemma 3.6 *The maximum number of topologically distinct perspective views of a polyhedral terrain with a total of n edges can be as large as $\Omega(n^8\alpha(n))$.*

Proof. The lower bound construction is similar to the construction of the previous subsection, and similar in spirit to the construction for perspective views of arbitrary polyhedra in [20]. We start with the same construction as for orthographic views, and for the extra degree of freedom that we now have, for moving the viewpoint closer to farther away from the scene, we use a displaced duplicate of construction B of the previous subsection, with $\Omega(n^3)$ changes in view as the viewpoint moves forwards or backwards.

Interpreted in terms of critical surfaces, the construction of the previous subsection induces two sets of surface patches: the surface patches of construction A are roughly parallel to the xy-plane, while the surface patches of construction B are roughly orthogonal to the first set, and are oriented say parallel to the xz-plane. (Recall that each surface induced by either construction is the collection of directions where three edges are coincident in the view, and two of the edges are arbitrarily close together. Thus, the resulting surface patches are almost flat.) We add the extra copy of construction B rotated 90° from the original construction, such that the critical surface patches induced by it will be roughly orthogonal to each of the other two sets, namely, they will be roughly parallel to the yz-plane. This results in a scene with $\Omega(n^8\alpha(n))$ different views. □

4 Arrangements of Convex Polyhedra

We next study arrangements defined by a collection of convex polyhedra in 3-space. These arrangements have the special property that a line in any direction stabs only a subset of the faces. Note that we consider the interior of each polyhedron as a portion of the arrangement. We prove the following theorem[6]

Theorem 4.1 *The maximum complexity of an arrangement induced by k convex polyhedra with a total of n vertices is $\Theta(nk^2)$.*

Proof. For the upper bound, note that the set S of k convex polyhedra has stabbing number $2k$ in any direction. Consider a segment $f \cap g$ for f, g faces of polyhedra in S. Then, either an edge e_f of f intersects g, or an edge e_g of g intersects f. By the stabbing property, each edge intersects at most $2k$ faces and hence there are at most $2nk$ segments $f \cap g$, over all faces f, g of all the polyhedra in S. Using the stabbing property once more, we see that each segment $f \cap g$ is intersected by at most $2k$ faces, and the upper bound follows.

To see that this bound is tight in the worst case, assume $k \leq n/3$ and take a convex polygon P_1 with n/k vertices[7] lying in the yz-plane. (If $k > n/3$ it is trivial to construct an arrangement with $\Omega(nk^2) = \Omega(n^3)$ vertices.) Denote the number of vertices of P_x by $|P_x|$. Duplicate P_1 $k/2 - 1$ times to obtain polygons $P_2, \ldots, P_{k/2}$ and rotate P_i slightly relative to P_{i-1} (see Figure 6). This results in a planar arrangement (in the yz plane) with complexity $\frac{1}{2}\sum_{i \neq j}(|P_i| + |P_j|) = \Omega(nk)$. Next, extend this arrangement in the x direction, and slice the resulting arrangement of cylinders with additional $k/2$ triangles, all parallel to the yz plane to get a subdivision of space with complexity $\Omega(nk^2)$. The overall number of polyhedron vertices in the construction is $n/2 + 3k/2 \leq n$. □

5 The Number of Views of Convex Polyhedra

In this section we study the number of views of a three-dimensional scene consisting of k non-intersecting opaque convex polyhedra having a total of n vertices. Again we consider two types of viewpoint space: The space related to orthographic views (from infinity) and the space related to perspective views. In [19] it was shown that for one convex polyhedron these bounds are $\Theta(n^2)$ and $\Theta(n^3)$ respectively. We first derive upper bounds on the maximum number of views of k convex polyhedra and then present lower bound constructions.

In Subsection 3.1 we have considered three different types of curves (corresponding to accidental viewpoints) that appear on the sphere of directions in the case of a polyhedral terrain. The same types of curves may occur in the case of convex polyhedra. As before, the curves of type III dominate the complexity of the arrangement, so we restrict our attention

[6] Independently, Aronov et al. [3] have obtained a similar result, generalized to arrangements of polytopes in d-dimensional space for a fixed d.

[7] In this proof m_1/m_2 should be interpreted as the integer $\lfloor m_1/m_2 \rfloor$.

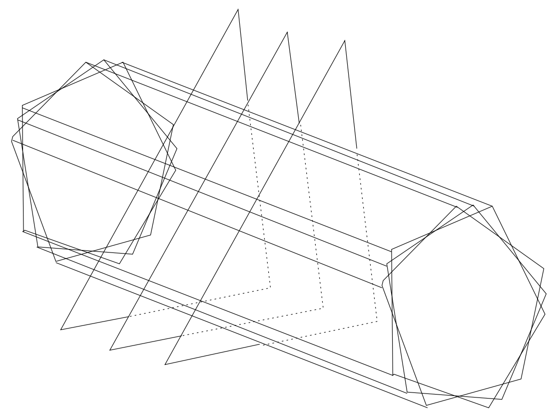

Figure 6: An arrangement of k convex polyhedra having $\Omega(nk^2)$ complexity

to them. Recall that a curve of type III represents a collection of viewing directions for which the views of a fixed triple of edges of the polyhedra meet at one point.

Let e_i, e_j and e_l be such a triple of edges. Assume for simplicity that each edge belongs to a distinct polyhedron P_i, P_j and P_l respectively. Each point on the curve represents a line L in the viewing direction that touches the three edges simultaneously. It is easy to verify that each such a line L is tangent to at least two of the polyhedra P_i, P_j, P_l. In other words, it may cross the interior of at most one of these polyhedra, as the polyhedra are opaque. Suppose that this is indeed the case and it crosses through P_l. Then, necessarily, the contact between L and P_l lies farther from the viewpoint than the contacts with P_i or P_j.

Next, we fix e_l and bound the possible number of curves of the third type, induced by e_l and pairs of edges—one edge of the polyhedron P_i and one of P_j. Denote the number of vertices of P_x by $|P_x|$.

Lemma 5.1 *The maximum number of pairs of edges, one from P_i and one from P_j, such that together with e_l they define a critical curve on the sphere of directions and such that e_l lies farthest from the viewpoint is $O(|P_i| + |P_j|)$.*

Proof. Suppose first that there is a plane Π_l that contains e_l such that both P_i and P_j lie on one side of Π_l. Take a plane Π parallel to Π_l, far away from the scene and such that e_l is farther from Π than P_i or P_j. Let $q = q(0)$ be an endpoint of e_l and draw on Π the intersection of all the lines through $q(0)$ that are tangent to P_i. The resulting curve on Π is evidently the boundary of a convex polygon, which we denote by $Q_i = Q_i(0)$. The polygon $Q_i(0)$ has at most $O(|P_i|)$ edges. As we let $q(t)$ move along e_l towards the other endpoint $q(1)$, Q_i will change continuously. Still, it will always remain a convex polygon. Furthermore, it will change its (combinatorial) structure only when the line through $q(t)$ coincides with a plane of a facet of P_i. Thus it will not have new edges appearing (or else have edges disappearing) more than $|P_i|$ times. The same arguments hold for P_j and its corresponding "shadow" Q_j on Π.

An intersection point of an edge of $Q_i(t)$ and an edge of $Q_j(t)$ along some interval $0 \leq t' < t < t'' \leq 1$ represents a curve of the third type on the sphere of directions. How many pairs of edges, one from each polygon, intersect on the boundary of the union of the two polygons? At $t = 0$ there are at most $(|Q_i(0)| + |Q_j(0)|) = O(|P_i| + |P_j|)$ such intersection. As t increases, every new pair of edges that intersect must be the result of a critical event that either makes the vertex of one polygon meet the edge of another, or that an edge of a polyhedron inducing a shadow edge is substituted by another edge of the same polyhedron. The first kind of critical event corresponds to the plane through a vertex of one polyhedron and the edge of the other polyhedron crossing e_l. For a fixed vertex of one polyhedron there are at most two edges of the other polyhedron that can participate in such an event, because we take the line through the vertex v and as we move it in contact with e_l it may be tangent to the other polyhedron, not containing v, at most twice. The second kind of critical event occurs when the line through $q(t)$ coincides with a plane of a facet of either polyhedron. This kind as well may incur at most two new intersections with the shadow of the other polyhedron. Therefore only $O(|P_i| + |P_j|)$ critical events occur as $Q_i(t)$ and $Q_j(t)$ move. And thus the overall number of potential curves involving e_l, P_i and P_j is at most $O(|P_i| + |P_j|)$.

To relax the assumption that there is a plane Π_l such that both polyhedra P_i, P_j lie on one side of it we do the following: We arbitrarily choose a plane Π_l containing e_l and cut each polyhedron that intersects Π_l by the plane Π_l into two. We repeat the analysis above for either side of Π_l and the polyhedra portions on that side. The only difference between the new situation and the previous one is that one or both of the corresponding Q_i and Q_j are now unbounded, but the entire analysis holds verbatim.

As to the assumption that the three edges e_i, e_j and e_l belong to three distinct polyhedra, one can easily verify that if the three edges belong to only one polyhedron or to two polyhedra, then the maximum overall number of critical curves that they can induce is asymptotically smaller than in the case where they belong to three distinct polyhedra. $\square$

Now we can state

Theorem 5.2 *The maximum number of topologically distinct views of a scene consisting of k convex non-intersecting polyhedra with a total of n vertices, when viewed from infinity, is $O(n^4k^2)$. The number of distinct views of such a scene where the viewpoint can be anywhere in space is $O(n^6k^3)$.*

Proof. Let $E(k, n)$ be the maximum number of curves of the third type that may appear on the sphere of directions in the current setting. Lemma 5.1 implies that for every edge e_l of a polyhedron P_l the number of critical curves of type III that it may induce due to interaction with a fixed pair of two additional polyhedra P_i and P_j is $O(|P_i| + |P_j|)$. Summing over all

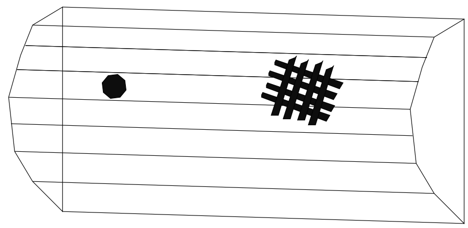

Figure 7: A construction giving $\Omega(n^2+nk^2)$ different views when the viewpoint moves along a meridian

edges e_l we get

$$E(k,n) \leq \sum_{i\neq j\neq l} |P_l| \cdot O(|P_i| + |P_j|) \leq 2\sum_{i=1}^{k}\sum_{j=1}^{k}\sum_{l=1}^{k} O(|P_j| \cdot |P_l|) = O(n^2k).$$

The two bounds of the theorem immediately follow. □

Finally, we exhibit lower bound constructions for the number of views of convex polyhedra.

Theorem 5.3 *A scene consisting of k convex polyhedra with a total of n vertices may induce $\Omega(n^4+n^2k^4)$ distinct views from infinity and $\Omega(n^6+n^3k^6)$ views when the viewpoint can be anywhere in space.*

Proof. Following the idea presented in Section 3 we first present a construction for one degree of freedom of the viewpoint that gives $\Omega(n^2+nk^2)$ different views when moving the viewpoint along a meridian on the sphere of directions. This construction is a superpositioning of two simpler constructions. The first consists of a "hill" with n horizontal edges in front of which there is a very small convex polygon with n edges. As we move the viewpoint up all the vertices of the polygon meet one edge of the hill in the view before they meet another edge (see Figure 7, on the left-hand side of the hill).

For the second construction we place a small slanted grid of roughly $k/2 \times k/2$ segments such that an edge of the hill meets all intersection points of the grid before another edge of the hill does so as the viewpoint moves up (see Figure 7, on the right-hand side of the hill). We position the polygon and the grid such that no two viewing events coincide as we move the viewpoint up or down.

We duplicate this construction and rotate the duplicate by 90° degrees to obtain a similar effect when moving from left to right. To obtain the bounds for perspective views we repeat the basic construction once again. □

6 Conclusion

In this paper we have shown an almost-tight combinatorial bound on the maximum number of topologically different orthographic views of polyhedral terrains. The bound is an order of magnitude lower than the corresponding bound for general polyhedra. We also analyzed the case of perspective views, and reduced this problem to the problem of bounding the complexity of a certain collection of segments defined relative to a terrain. We obtained these results by investigating arrangements of objects (curves or surfaces) that have the special property that every vertical line stabs only a small number of the objects. We believe that our results for this type of arrangements are of independent interest. We also presented extensions of these results to higher dimensions. Furthermore, we have presented bounds on the number of views of a scene consisting of k convex polyhedra with a total of n vertices.

We suggest the following open problems:

1. Tighten the gap between the lower and upper bounds on the number of views of k convex polyhedra with n vertices in total. A possible approach to improve the upper bound would be to obtain a low stabbing number in the spirit of the result by Cole and Sharir stated as Theorem 3.1 here.

2. What is the complexity of arrangements of surfaces that have a low stabbing number in more than one direction? For example, it would be interesting to have such a bound as a function of n, k_{min} and k_{max}, where k_{min} and k_{max} are the minimum and maximum stabbing number in any direction.

3. What is the maximum complexity $\nu(n)$ of the special set of segments lying above a terrain as defined in Subsection 3.2?

Our paper has concentrated on the combinatorial questions concerning aspect graphs of certain polyhedral scenes. We have not addressed the related algorithmic issues. Efficient computation of a sparse 2D arrangement is straightforward using *plane sweep* (see, e.g., [21]). We believe that computing a sparse 3D arrangement of surfaces in time that is roughly proportional to the maximum combinatorial complexity of the arrangement is fairly simple, imitating the proof of Proposition 2.3, although there are several technical details that still need to be studied. A somewhat more challenging problem is to compute an arrangement of convex polyhedra efficiently.

Acknowledgement

The authors thank Pankaj Agarwal for pointing out an error in Section 3 in an earlier draft of the paper. We also thank two anonymous referees for many helpful comments.

References

[1] P. K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Discrete and Computational Geometry* **12** (1994), pp. 177–182.

[2] P. K. Agarwal, M. Sharir and P. W. Shor, Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences, *J. Combinatorial Theory, Series A* **52** (1989), pp. 228–274.

[3] B. Aronov, M. Bern and D. Eppstein, Arrangements of polytopes, *manuscript*, 1991.

[4] K. W. Bowyer and C. R. Dyer, Aspect graphs: An introduction and survey of recent results, *Int. J. of Imaging Systems and Technology* **2** (1990), pp. 315–328.

[5] G. Castore, Solid modeling, aspect graphs and robot vision, in *Solid Modeling by Computer* (Pickett and Boyse, Eds.), Plenum Press, New-York, 1984, pp. 277–292.

[6] I. Chakravarty and H. Freeman, Characteristic views as a basis for three-dimensional object recognition, *Proc. SPIE: Robot Vision*, Vol. 336 (1982), pp. 37–45.

[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 ICALP*, 1989, pp. 179–192. Also in *Theoretical Computer Science* **84** (1991), pp. 77–105.

[8] K. L. Clarkson and P. W. Shor, Applications of random sampling in computational geometry, II, *Discrete and Computational Geometry* **4** (1989), pp. 387–421.

[9] R. Cole and M. Sharir, Visibility problems for polyhedral terrains, *Journal of Symbolic Computation* **7** (1989), pp. 11–30.

[10] H. Edelsbrunner, *Algorithms in Combinatorial Geometry*, Springer-Verlag, Heidelberg, 1987.

[11] Z. Gigus, J. Canny and R. Seidel, Efficiently computing and representing aspect graphs of polyhedral objects, *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13** (1991), pp. 542–551.

[12] L. Guibas and M. Sharir, Combinatorics and algorithms of arrangements, in *New Trends in Discrete and Computational Geometry* (J. Pach, Ed.), Springer-Verlag, 1993, pp. 9–36.

[13] 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.

[14] M. Hebert and T. Kanade, The 3D profile method for object recognition, *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, 1985, pp. 458–463.

[15] J. J. Koenderink and A. J. van Doorn, The singularities of visual mapping, *Biological Cybernetics* **24** (1976), pp. 51–59.

[16] J. J. Koenderink and A. J. van Doorn, The internal representation of solid shape with respect to vision, *Biological Cybernetics* **32** (1979), pp. 211–216.

[17] D. J. Kriegman and J. Ponce, Computing exact aspect graphs of curved objects: Solids of revolution, *International Journal of Computer Vision* **5** (1990), pp. 119–135.

[18] J. Matoušek, Approximations and optimal geometric divide-and-conquer, in *Proc. 23rd ACM Symp. Theory of Comp.*, 1991, pp. 505–511.

[19] W. H. Plantinga and C. R. Dyer, An algorithm for constructing the aspect graph, in *Proc. 27th IEEE Symp. Foundations of Computer Science*, 1986, pp. 123–131.

[20] W. H. Plantinga and C. R. Dyer, Visibility, occlusion and the aspect graph, *Int. J. of Computer Vision* **5** (1990), pp. 137–160.

[21] F. P. Preparata and M. I. Shamos, *Computational Geometry—An Introduction*, Springer-Verlag, New York, 1985.

[22] M. Sharir, On k-sets in arrangements of curves and surfaces, *Discrete and Computational Geometry* **6** (1991), pp. 593–613.

[23] J. Snoeyink, The number of views of axis-parallel objects, *Algorithms Review* **2** (1991), pp. 27–32.

[24] T. Sripradisvarakul and R. Jain, Generating aspect graphs for curved objects, in *Proc. IEEE Workshop on Interpretation of 3D scenes*, New York, 1989, pp. 109–115.

[25] J. Stewman and K. Bowyer, Creating the perspective projection aspect graph of polyhedral objects, *Proc. IEEE International Conference on Computer Vision*, New York, 1988, pp. 494–500.

[26] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport-Schinzel sequences by segments, *Discrete and Computational Geometry* **3** (1988), pp. 15–47.