Available online at www.sciencedirect.com
ScienceDirect
Computational Geometry 38 (2007) 129–138
Computational Geometry
Theory and Applications
www.elsevier.com/locate/comgeo

On incremental rendering of silhouette maps of a polyhedral scene ☆

Alon Efrat [a,*], Leonidas J. Guibas [b,1], Olaf A. Hall-Holt [c], Li Zhang [d]

[a] *The University of Arizona, Department of Computer Science, Gould-Simpson Building, PO Box 210077, Tucson, AZ 85721-0077, USA*
[b] *Computer Science Department, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA*
[c] *Department of Mathematics, Statistics, and Computer Science St. Olaf College, USA*
[d] *Hewlett Packard Labs, Mail Stop 1250, 1501 Page Mill Road, Palo Alto, CA 94304, USA*

Received 8 May 2003; received in revised form 25 March 2006; accepted 19 December 2006

Available online 23 January 2007

Communicated by R. Tamassia

Abstract

We consider the problem of incrementally rendering a polyhedral scene while the viewpoint is moving. In practical situations the number of geometric primitives to be rendered can be as large as many millions. It is sometimes advantageous to render only the silhouettes of the objects, rather than the objects themselves. Such an approach is regularly used for example in the domain of *non-photorealistic rendering*, where the rendering of silhouette edges plays a key role. The difficult part in efficiently implementing a kinetic approach to this problem is to realize when the rendered silhouette undergoes a *combinatorial change*.

In this paper, we obtain bounds on several problems involving the number of these events for a collection of k objects, with a total of n edges. We assume that our objects are convex polytopes, and that the viewpoint is moving along a straight line, or along an algebraic curve of bounded low degree. We also study the special case when the scene is a *polyhedral terrain*, and present improved bounds for this case. In addition to bounding the number events, we also obtain algorithms that compute all the changes occurring during a linear motion both for general scenes and for terrains.

Keywords: Incremental rendering; Visibility graphs; Silhouette arrangements

1. Introduction

There is an increasing demand in computer graphics applications for rendering large and complex environments involving scenes with many millions of polygons. The computational demands of such a task have to be addressed by both improved hardware and better algorithms. The very high complexity of these environments in terms of simple

☆ A preliminary version of this paper appeared in [A. Efrat, L.J. Guibas, O.A. Hall-Holt, L. Zhang, On incremental rendering of silhouette maps of a polyhedral scene, in: Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2000, pp. 910–917].

* Corresponding author.

E-mail addresses: alon@cs.arizona.edu (A. Efrat), guibas@cs.stanford.edu (L.J. Guibas), olaf@stolaf.edu (O.A. Hall-Holt), l.zhang@hp.com (L. Zhang).

URLs: http://www.cs.arizona.edu/people/alon/ (A. Alon), http://graphics.stanford.edu/~guibas (L.J. Guibas), http://www.stolaf.edu/people/olaf (O.A. Hall-Holt), http://www.hpl.hp.com/personal/lizhang (L. Zhang).

1 Partially supported by NSF grant CCR-9623851 and by ARO MURI grant DAAH04-96-1-007.

0925-7721/$ – see front matter © 2007 Published by Elsevier B.V.
doi:10.1016/j.comgeo.2006.12.003

geometric primitives, such as triangles, is in part an artifact of the traditional rendering pipeline of current graphics systems, which are based on triangle scan-conversion as the basic primitive. In general, the number of different objects present in a scene is much less than millions—and the high triangle count is due to the tessellation of more complex curved objects into polyhedral approximations that can be rendered by the hardware.

Triangle edges have to be handled properly in order to obtain high quality renderings of a scene. There is a vast literature in computer graphics on how to deal with edge problems such as jaggies, antialiasing, etc. Yet it is important to realize that not all edges are created equal. Edges in the rendered image separating two different objects are much more likely to be problematic than edges separating two polygons belonging to the same object. Across the former we will have depth discontinuities, a different reflectance function on the two sides, different colors, etc. Across the latter simple edges, interpolatory smoothing techniques work well to simulate the appearance of a smooth surface. The former edges are *silhouette* edges. We study in this paper combinatorial properties and algorithms related to silhouette edges.

We consider a small number of objects that have been tessellated into a much larger number of triangles. Given a viewpoint, each object has a *silhouette*, a collection of edges forming closed cycles that separate triangles visible from triangles invisible to the viewpoint. We focus on the geometric structure of these silhouettes and how they subdivide the plane (their arrangements). Such silhouette structures are important in efficient rendering. For example, we can calculate the shadow and perform occlusion culling by using the silhouettes of the objects, or we can answer ray shooting queries efficiently if we know silhouette structures. In recent work [13], Gu et al. present an efficient image-based rendering method by using the object silhouette to clip the images rendered on a coarse geometry. If we can compute these silhouettes for the viewpoint, and also maintain them as the viewpoint moves around (incremental rendering), not only do we know the most important edges in the image to be rendered, but we can also facilitate many other rendering operations.

The difficult part in efficiently implementing any algorithm for rendering moving objects is to realize when the image undergoes a *combinatorial change*, defined as a change where either the polytope edges that define the silhouette changes or the topology of the silhouette structure changes e.g., when two silhouettes start or stop intersecting or when a hole in their union appears or disappears.

The input to the first type of problems we investigate is a set $\mathcal{S} = \{P_1 \dots P_k\}$ of convex polytopes in 3-D that we have to render, and a viewpoint p. As mentioned, we make the realistic assumption that the number of polytopes k is much smaller than the total number n of vertices of these polytopes. Let S be a plane representing the display surface. The *shadow* of an object is the perspective projection of the object on S from p. The *silhouette* of a polytope is the boundary of its shadow, which is a convex polygon. The different silhouettes subdivide S into regions, so that one cannot move in S from one region into another without crossing at least of one the silhouettes. This structure is called the *silhouette arrangement*. We refer the reader to [6] for more details about arrangements. The *silhouette map* is the arrangement on S with the hidden part removed. Formally, we assign a unique color (or ID) to each object. For any point q on the background sphere S, assign to q the color of the first object hit by the ray starting from the perspective point and shooting toward q. The boundary of the monochromatic regions is exactly the silhouette map. The *silhouettes-of-union* (or *silhouettes-union* for short) is the boundary of the union of all the shadows on S.

We describe combinatorial bounds on these geometric structures, in each of the following three cases

- *Static viewpoint*—the viewpoint p is static.
- *Linear motion*—p moves along a straight line, and the goal is to bound the number of combinatorial changes each of the structures undergoes.
- *Algebraic motion*—p moves along an algebraic curve, and again the goal is to bound the number of changes each of the structures goes through.

In this paper we present these and other bounds using the following notation. Let $\beta_s(n) = \lambda_s(n)/n$, where $\lambda_s(n)$ is the maximum length of a Davenport–Schinzel sequence of order s with n symbols; see [21]. When s is constant, $\beta_s(n)$ is an extremely slowly growing function—in particular, $\beta_3(n) = \alpha(n)$, the functional inverse of Ackermann's function.

The bounds are summarized in Table 1. For the static viewpoint, the bounds given are for the combinatorial size of the respective structure. For the moving viewpoint columns, the bounds indicate the number of combinatorial changes in the relevant structure. For the lower bound of the number of silhouettes-unions changes under the linear motion, we

Table 1
The bounds on the complexity of silhouette structures

Structure	Static viewpoint	Linear motion	Algebraic motion
silhouette arrang.	$\Theta(kn)$	$\Theta(k^2 n)$	$\Theta(kn^2)$
silhouette map	$\Theta(kn)$	$\Theta(k^2 n)$	$\Theta(kn^2)$
union of silhouette	$\Theta(n\alpha(k) + k^2)$	$\Theta(k^2 n)$	$O(kn^2)$ $\Omega(n^2\alpha(k) + k^2 n)$

proved a weaker bound of $\Omega(kn\alpha(k) + k^3)$ in the conference version of the paper [11], and the tight lower bound of $\Omega(k^2 n)$ was obtained later by Aronov et al. [2].

For terrains, we consider the measure in terms of the number of 'mountains' in a terrain. Roughly speaking, a mountain is a up-convex body with its base on the xy-plane. For a terrain with k mountains and n vertices and a vertically moving viewpoint, we are able to obtain a roughly $\Theta(kn)$ bound on the number of combinatorial changes of the silhouette map and silhouettes-union. When a terrain has n vertices, we can always decompose it into $O(n)$ mountains where each mountain is formed by taking the vertical prism under a triangle. Although in the worst case, k can be as large as $\Theta(n)$, in real applications, the number of mountains is usually small. This bound is then better than the roughly $\Theta(n^2)$ bound on the number of combinatorial changes of the silhouettes map for a vertically moving viewpoint where the lower bound can be achieved by two mountains.

One key lemma we use to achieve the upper bounds is to bound the number of so called EEE events, i.e. the number of times when the shadow of three edges on the silhouettes come together. This quantity is in turn obtained by bounding the number of lines that touch three convex polytopes and pass through the perspective point during the entire motion. We show that, for a linearly moving point, this number is linear in the total complexity of these three convex polytopes instead of quadratic, and thus obtain better upper bounds.

Based on the upper bound on the combinatorial changes, we can obtain algorithms that compute all the changes occurring during a linear motion, both for general scenes and for terrains in time $O(k^2 n \log n)$ and $O(kn\alpha(n)\log^2 n)$, respectively.

Related results. Similar problems were investigated both analytically (usually in the Computational Geometry community) and empirically (in the Computer Graphics community). Among the analytic results, several authors considered the problem of computing the visibility maps from a moving viewpoint [7,16,18]. In [5], de Berg, Halperin, Overmars, and van Kreveld presented a number of results regarding the complexity of the aspect graph for different scenarios, and its relation to the complexity of arrangements. Barequet et al. [3,19] showed how to use the BAR-tree to obtain fast rendering of the silhouette of a (not necessarily convex) polytope. Lower bounds on the number of different viewpoints and silhouette maps for arbitrary location of the viewpoint were obtained by Aronov et al. [2], that also describes a lower bound of $\Omega(k^2 n)$ for the union of silhouette changes under the linear motion.

Kettner and Welzl [15] studied the complexity of the approximated silhouette of a convex and non-convex scene, where the approximation is considered under the Hausdorff metric, in the projection on the unit sphere.

Other works: [8,12–14,17,20] in the graphics community also use shadows and silhouettes as a means to simplify the description of a complicated environment. Silhouettes are also useful in collision detection [9], and other applications.

2. Lower bounds

2.1. Silhouette structures from a static point

From a fixed viewpoint, the shadow of each polytope is a convex polygon on the background plane. For k convex polygons with n vertices in total, tight bounds of $\Theta(kn)$ and $\Theta(n\alpha(k) + k^2)$ are known for, respectively, the complexity of their arrangement and the boundary of their union [1]. These bounds yield the tight bounds for silhouette arrangements and silhouettes-union. For the silhouette map, we will present an example to show complexity of $\Omega(kn)$ and thus obtain a tight bound of $\Theta(kn)$ for silhouette maps.

Fig. 1(left) illustrates an example of our construction of k fat convex polygons, and Fig. 1(right) shows an enlargement of a neighborhood containing k polygon corners. The goal of the construction is to create n/k corner

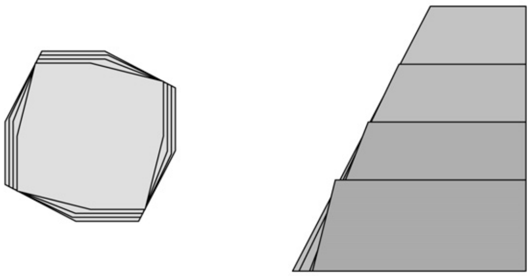

Fig. 1. The lower bound construction of the silhouette map. Left: The global arrangement. Right: Close up of a corner.

neighborhoods with silhouette map complexity $\Omega(k^2)$ each, yielding a total complexity of $(n/k) \times \Omega(k^2) = \Omega(kn)$. The corner of a polygon consists of the corner vertex, and two incident edges, which we will call the left and right edges. Let us focus on a particular corner neighborhood N (such as the one depicted in Fig. 1(right)). To obtain the desired complexity, we need to ensure that the left edge of each polygon corner contributes a vertex to the silhouette map at the right edge of all preceding polygons in the depth order. Intuitively, we seek a set of nearly parallel left edges with increasing slope, such that each left edge lies to the right of the endpoints of the previous left edges (in the depth order).

To verify that this is possible, consider a sequence of tangents to the unit hyperbola in the first quadrant of the plane ($y = \sqrt{x^2 - 1}$). Let the tangent points be a_i, where they are given in order of increasing x-coordinate. For a tangent to the hyperbola at a_i, let b_i be the intersection of the tangent with the asymptote $y = x$. Let the b_i be our corner vertices, and let the left edge from b_i extend down toward a_i. Thus a left edge must pass below the corner vertices of all preceding polygon corners. We still need to specify the spacing of the points a_i. First let all the right edges proceed to the right, with slope $0 < \alpha < 1$. Now let a_{i+1} be the intersection of the previous right edge with the hyperbola. That is, we make each right edge go up toward the hyperbola, and choose all tangent points except the first one to be the intersections of the right edges with the hyperbola. Thus we also ensure that any left edge will be above the tangent points corresponding to all previous left edges, and hence visibly intersecting all previous right edges.

To construct the k polygons from such a neighborhood, we place n/k rotated copies of the neighborhood near the vertices of a sufficiently large $\frac{n}{2k}$-gon, and connect corresponding left and right edges. In order for this to work, we need to make sure that the corner angles are sufficiently large. In the neighborhood construction the position of the first tangent point and the parameter α were left unspecified. These together determine a lower bound on all the corner angles, which can be anything less than π, so we can always make them large enough.

In summary, we have that:

Theorem 2.1. *In the worst case, the silhouette arrangement, silhouette map, and silhouettes-union have complexity* $\Theta(kn)$, $\Theta(kn)$, *and* $\Theta(n\alpha(k) + k^2)$, *respectively.*

2.2. Lower bounds for a moving perspective viewpoint

Now that we have tight bounds for a fixed viewpoint, we consider the number of changes for a moving viewpoint. First, we will give a generic construction for lower bounds. Suppose that σ is the maximum complexity from a static perspective point p. There exists a small ball B around p such that for any point q inside B, the complexity of the structure from q has complexity σ. For linear motion, we can put k line segments inside B so that when we move the viewpoint p linearly, the shadow of those k line segments sweep over the silhouette structure. This way, we create $k\sigma$ changes. For algebraic motion, we take the classical example of a quadric curve intersecting a convex n-gon n times. Then, in this manner, we can create $n\sigma$ changes for algebraic motion. Combining this approach with the maximum possible complexity for a static point, we have the following:

Theorem 2.2. *For linear motion, the silhouette arrangement, silhouette map, silhouettes-union can change* $\Omega(k^2 n)$, $\Omega(k^2 n)$, $\Omega(kn\alpha(k) + k^3)$ *times, respectively. The bound on the first two quantities is tight, and a better* (*and tight*) *lower bound of* $\Omega(k^2 n)$ *on the third quantity was obtained in* [2]. *For algebraic motion, the lower bounds are* $\Omega(kn^2)$, $\Omega(kn^2)$, *and* $\Omega(n^2\alpha(k) + k^2 n)$, *respectively.*

3. Upper bounds

In the paper of de Berg et al. [5], they bound the number of different views to a scene of k convex objects. They derived an upper bound of $O(kn^2)$ on the number of surface patches that form a partition of the viewpoint space into cells with the same combinatorial structure of the visibility. Since a constant degree algebraic curve can intersect such a surface patch only a constant number of times, we can obtain the following upper bound by borrowing Theorem 5.2 of [5].

Theorem 3.1. *For a scene that consists of k convex polyhedra with total complexity n, there are $O(kn^2)$ combinatorial changes to the silhouette arrangement when the perspective point moves along a constant degree algebraic curve.*

However, when the point moves linearly, we obtain better bounds on the number of changes of the silhouette arrangement.

Theorem 3.2. *For a scene that consists of k convex polyhedra in general position with total complexity n, there are $O(k^2n)$ combinatorial changes to the silhouette arrangement when the perspective point moves linearly. This bound also applies to the number of changes in the silhouette-map.*

For a convex polyhedron P, a line ℓ is said to be *tangent* to P if ℓ meets P but not its interior. When ℓ intersects P at exactly one point, it is called *strictly tangent* to P. The following simple fact is useful.

Fact 3.3. *If ℓ is strictly tangent to P, it is strictly tangent to $P \cap \gamma$ for every plane γ containing ℓ. Conversely, if ℓ is strictly tangent to $P \cap \gamma$ for some plane γ containing ℓ, it is strictly tangent to P.*

Lemma 3.4 is the key step to prove Theorem 3.2. It shows that the number of lines that touch a given line and three convex polytopes is linear in the total complexity of those polytopes. We believe that this lemma is of interest as an independent result.

Lemma 3.4. *For any given line ℓ and three convex polyhedra P_1, P_2, and P_3 in general position, the number of lines that touch ℓ and are tangent to P_1, P_2, and P_3 is $O(|P_1| + |P_2| + |P_3|)$. Here $|P_i|$ is the number of vertices of P_i (for $i = 1, 2, 3$).*

In what follows, without loss of generality, we assume that the line ℓ is the z-axis. Consider the family Γ of planes that pass through the z-axis. We parameterize them according to the angles they make with the x-axis: $\Gamma = \{\gamma(\theta)\colon 0 \leqslant \theta < \pi\}$. For a convex polyhedron P, denote by $P(\theta)$ the intersection between P and $\gamma(\theta)$. Clearly, $P(\theta)$, if not empty, is a convex polygon lying on $\gamma(\theta)$. For two convex polyhedra P, Q and any θ, $0 \leqslant \theta < \pi$ where $P(\theta)$ and $Q(\theta)$ are not empty, we define the *lower outer bi-tangent* between $P(\theta)$ and $Q(\theta)$ to be the line which is the common tangent to $P(\theta)$ and $Q(\theta)$, and has $P(\theta)$, $Q(\theta)$ and the point $z = \infty$ on the same side. Let $\phi_{P,Q}(\theta)$ (or $\phi(\theta)$ if P, Q are clear from the context) denote its slope. Let $\phi_{P,Q}$ denote the graph of the function $\phi_{P,Q}(\theta)$ as θ varies.

First, we observe that:

Lemma 3.5. *The graph of the function $\phi_{P,Q}(\theta)$ consists of $O(|P| + |Q|)$ arcs, each of which is a constant-degree rational function of $\tan(\theta)$.*

Proof. For any particular θ, the cross intersection $P(\theta)$ is a convex polygon. A vertex of this polygon is either a vertex v of P or $e \cap \gamma(\theta)$ for an edge e of P. For all the vertices created by the same edge, we think of them as a single vertex moving on a low degree rational curve as the plane rotates. As long as a bi-tangent is defined by the same pair of vertices, its slope is just a rational function in terms of $\tan\theta$. When can a breakpoint be created? There are two possibilities: first, when either a previous vertex of the cross section is deleted or a new one is created; second, when three vertices, two from one polygon and one from the other, are collinear and the line that passes through them is a bi-tangent line. Clearly the first type of event can happen at most $O(|P| + |Q|)$ times as at this event the plane sweeps over a vertex of P or of Q.

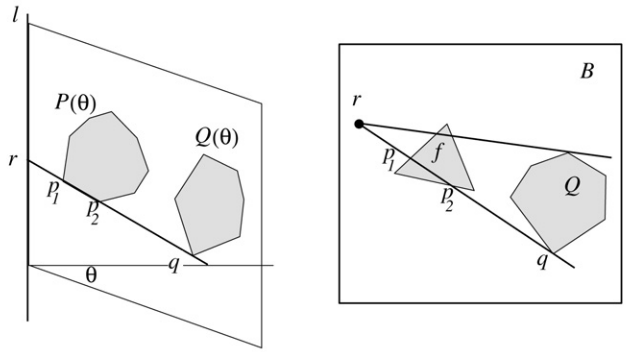

Fig. 2. Combinatorial change to a bi-tangent.

For the second type of event, suppose that for some θ, $p_1, p_2 \in P(\theta)$ and $q \in Q(\theta)$ are collinear and the line ℓ' determined by $p_1 p_2 q$ is a bi-tangent line to $P(\theta)$, $Q(\theta)$. First of all, p_1, p_2 must be adjacent vertices on $P(\theta)$ by convexity. Further, by the general position assumption, ℓ' is strictly tangent to $Q(\theta)$ at q, which implies that ℓ' is strictly tangent to Q according to Fact 3.3.

Consider the edges $e_1, e_2 \in P$ that correspond to p_1, p_2, i.e. $p_1 = e_1 \cap \gamma(\theta)$ and $p_2 = e_2 \cap \gamma(\theta)$. We know that e_1, e_2 must lie on the same face of P, say f (Fig. 2 right). Consider the plane B that supports f. Observe that $B \cap Q$ is again a convex polygon. Suppose that B intersects ℓ at point r. We claim that rq is strictly tangent to $B \cap Q$. This simply follows from the fact that ℓ' is indeed strictly tangent to Q and by Fact 3.3. From any point, we can draw at most two tangent lines to another convex polygon. This is to say that for any face f of P, there are at most two such points q on Q such that for some $p_1, p_2 \in f \cap \gamma(\theta)$, q is collinear with p_1, p_2 as a bi-tangent line. Therefore, the second type of event can happen at most $O(|P| + |Q|)$ times.

This concludes the proof of the lemma. □

Building on Lemma 3.5, we now can prove Lemma 3.4.

Proof of Lemma 3.4. For P_1, P_2, P_3, we plot two functions $\phi_1(\theta) = \phi_{P_1,P_2}(\theta)$ and $\phi_2(\theta) = \phi_{P_1,P_3}(\theta)$. By Lemma 3.5, $\phi_1(\theta)$ consists of $O(|P_1| + |P_2|)$ low-degree rational arcs and $\phi_2(\theta)$ consists of $O(|P_1| + |P_3|)$ such arcs. The graphs ϕ_1 and ϕ_2 can intersect at $O(|P_1| + |P_2| + |P_3|)$ points. For a line ℓ' that touches ℓ, consider the plane γ that is determined by ℓ and ℓ'. A line ℓ is tangent to P_1, P_2, P_3 if it is a common tangent to $P_1 \cap \gamma$, $P_2 \cap \gamma$, and $P_3 \cap \gamma$. This must correspond to an intersection point between ϕ_1 and ϕ_2. Therefore, the total number of such lines is bounded by $O(|P_1| + |P_2| + |P_3|)$. Of course, there are different types of tangents. But this is not a problem as there are a constant number of different combinations of tangent types (16 combinations to be precise). We have therefore proved Lemma 3.4. □

Next, we proceed to prove Theorem 3.2.

Proof of Theorem 3.2. The silhouette arrangement changes combinatorially only when the edges that define a silhouette change, or when the silhouettes of two polytopes start or stop intersecting, or when three silhouette edges intersect at a single point. These changes correspond to the following three types of events.

1. The *first type* (F) occurs when the viewpoint crosses a plane supporting a facet.
2. The *second type* (VE) occurs when the viewpoint crosses a plane determined by a vertex and an edge from a different polytope.
3. The *third type* (EEE) occurs when there is a ray from the viewpoint that touches three different polytopes.

The number of the events of the first type is bounded by $O(n)$, as there are $O(n)$ facets in total. At an event of this type an edge of the silhouette is replaced by two other edges, or two edges become collinear and are replaced by a single edge. Events of this type causes at most O(k) changes to the silhouette arrangement. Therefore, this type of event causes O(kn) changes altogether.

For an event of type two, suppose the vertex and polytope involved are v and P, respectively. Consider the double cone C formed by the union of all lines passing through v and P. The event happens only when the viewpoint on ℓ crosses the boundary of C, which can happen at most twice. This implies that the second type of event can happen at most O(kn) times. Once such event happens, it can cause O(1) changes to the silhouette as it makes a vertex cross an edge in the silhouette.

The third type is the difficult case, when there is a line from the viewpoint that goes through the boundary of three polytopes. By Lemma 3.4, we know that the number of times this happens is bounded by

$$\sum_{i,j,k} \mathrm{O}\big(|P_i| + |P_j| + |P_k|\big) = \mathrm{O}(k^2 n),$$

and each causes a single change in the silhouette arrangement.

This concludes the proof of the theorem. □

Notice that we actually upper-bounded the number of changes of the silhouette arrangement, and therefore the silhouette map and silhouettes-union. For the silhouette arrangement and silhouette map, these upper bounds match the lower bounds of Theorem 2.2. For the silhouettes-union, a matching lower bound was shown by Aronov et al. [2].

4. Terrain with k mountains

As another application of Lemma 3.5, we bound the number of changes of the silhouette for our special terrains, namely, terrains that consist of mountains. A convex object M is called a *mountain* if for any vertical line ℓ, $\ell \cap M$ (if not empty) has one endpoint on the xy-plane, and one endpoint above the xy-plane. Intuitively, a mountain is an upper-convex object whose base is on the xy-plane.

In the case of visibility, Cole and Sharir showed in [10] that for a viewpoint moving vertically in a terrain with n vertices, the visibility changes O($n^2 2^{\alpha(n)}$) times, improving the naive bound of O(n^3). Having a small number of mountains does not help to reduce the lower bound there as the $\Omega(n^2)$ lower bound can be constructed by using *two* mountains. On the other hand, consider the silhouette arrangement from a vertically moving viewpoint. It is not difficult to construct an example similar to Fig. 1 to show that there can be $\Omega(k^2 n)$ changes to it, matching the upper bound given in Theorem 3.2. However, we obtain a better bound on the number of combinatorial changes of the silhouette map (and silhouettes-union) for a viewpoint moving vertically.

Theorem 4.1. *The silhouette map and silhouettes-union change* $\Omega(kn)$ *and* O($kn\beta_4(k)$) *times for a point moving vertically in a terrain with k mountains and n vertices in the worst case.*

Proof. For the lower bound, consider the picture in which we have a cylindrical mountain P with n sides (see Fig. 3). In front of the mountain, we have another k peaks (skinny tetrahedra). Then when the viewpoint moves vertically, each time it crosses a plane supporting a facet of P, it causes $\Omega(k)$ changes to the silhouettes-union and thus the silhouette map. In total, the number of changes is $\Omega(kn)$.

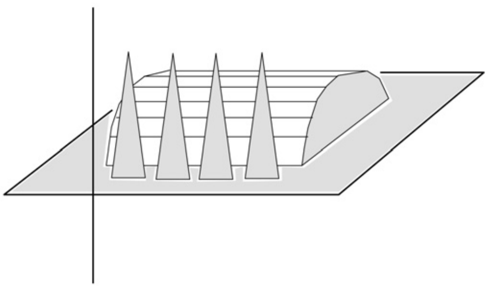

Fig. 3. Lower bound construction for terrain.

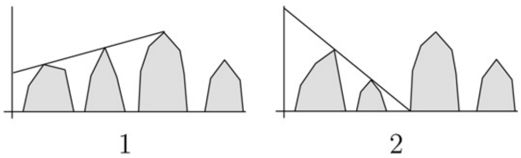

Fig. 4. Lines that touch z-axis and three mountains.

For the upper bound, consider again the three types of event used previously. The first two cases are bounded by $O(kn)$ as we have seen from the argument for general convex polytopes.

To bound the events of the third type, we now need to bound the number of lines that touch the z-axis and three mountains and avoid all the other mountains. Again consider Γ, the family of the planes that pass through the z-axis. There are two cases where such a line can appear. One case is when there is a line that is tangent to three mountains from the same side and avoids all the other mountains; and the other case is when there is a line that touches the base (intersection of the mountain with the xy-plane) of one mountain and is tangent to two other mountains. See Fig. 4.

As in [10], we define a partial ordering on the mountains so that $P \prec Q$ if there is a ray emanating from a point on the z-axis that intersects both P, Q and hits P first. Since the P_i's are mountains, this ordering is consistent. Among the three polyhedra involved in an event, we call the object that is furthest from the z-axis (i.e. the greatest one under the partial ordering $\prec$) the *dominant* object.

Now, let us focus on one polytope, say P_1. We will bound the number of events in which P_1 is the dominant object. For the first case, such an event happens when there is a line passing through the z-axis and tangent to P_i, P_j, P_1 from above, where $P_i \prec P_j \prec P_1$. In addition, this line must not intersect any other object P for $P \prec P_1$.

For an object P_j for which $P_j \prec P_1$, define the function $\phi_j(\theta)$ to be the slope of the upper outer bi-tangent of $P_1(\theta)$ and $P_j(\theta)$, if both $P_1(\theta)$, $P_j(\theta)$ are non-empty. Otherwise, $\phi_j(\theta)$ is undefined. Then, the above event happens when $\phi_i(\theta), \phi_j(\theta)$ have the same value and they are the highest among all the $\phi_l(\theta)$'s for $P_l \prec P_1$. In other words, we can charge each of these events to a unique break point on the upper envelope of the ϕ_j's. As we have shown in Lemma 3.5, the functions $\phi_j(\theta)$ consist of $O(|P_1| + |P_j|)$ algebraic arcs. Further, for any two algebraic arcs, they can intersect at most twice, as there can be at most two lines intersecting four lines in general position in 3-spaces. Applying the standard argument of cutting the curves into intervals with each interval containing $O(k)$ arcs, that this number of events is bounded by $\beta_4(k) \sum_j (|P_1| + |P_j|)$. Summing this up for all the polytopes, we have:

$$\sum_i \left(\beta_4(k) \sum_j \big(|P_i| + |P_j|\big)\right) = kn\beta_4(k).$$

Similarly, we can bound the number of events in the second case by defining another function $\xi_j(\theta)$ as the slope of the line incident to the base of P_1 and tangent to P_j from above.

To summarize, the number of changes of the silhouette for a vertically moving point is $O(kn\beta_s(k))$ in a terrain with k mountains and n vertices. □

5. Algorithms

We can apply the above combinatorial bounds for a linearly moving viewpoint, both in a general scene and for our terrains, to devise algorithms to maintain the different silhouette structures for a viewpoint that moves along an input line. Assume that the line along which the viewpoint moves is the z-axis. Given a scene with k convex objects, we first compute the critical points on the line at which the silhouette structure changes. For F-events, we can simply compute them for each facet. It would be too expensive if we use the bruteforce method for VE and EEE events. We describe a sweeping method with a lower complexity.

Image a plane γ that rotates around the z-axis from 0 to π. Consider the intersection of γ and the convex polytopes, which is a set of convex polygons. When γ rotates, those polygons deform and move. We can easily maintain the combinatorial structure of each polygon by tracking when γ passes a polytope vertex. To detect VE events, each time when γ sweeps over a vertex v, we need to compute the tangent from v to each polygon. This can be done in $O(k \log n)$ time by a binary search against each polygon. In total, it takes $O(kn \log n)$ time to compute critical points associated with VE events. For EEE events, we wish to detect when a common tangent line to three convex polygons arises during the motion. We first track the bi-tangents of each pair of convex polygons by checking locally when the bi-tangent becomes collinear with an adjacent edge. Then we maintain a list of all the bi-tangents sorted according to their slopes. The task of detecting when an event happens reduces to detecting when the slopes of two bi-tangent lines coincide. This way, we can detect all the lines that touch three convex polyhedra and the z-axis. It can be seen that the events that happen in our algorithm can be counted exactly as those counted by Theorem 3.2 and Lemma 3.5. The working space is $O(k^2 + n)$ and the processing time for each event is $O(\log n)$.

We can apply similar algorithms to terrains. The only difference is that instead of maintaining the sorted list of all the tangents to an object, we maintain the lowest (or highest) one according to which side it is on, as described in the proof of Theorem 4.1. This can be done by the kinetic tournament data structure presented in [4].

During the process of sweeping, we can also record the event associated with each critical point. Those events can be used to update the silhouette arrangement when the view point moves. The update time, however, varies according to the type of events. For F even, it may cause $\Theta(k)$ changes and takes $\Theta(k)$ time to update the silhouette. For VE and EEE events, the update to the silhouette map takes $O(1)$ time.

Thus, we have:

Theorem 5.1. *For k convex polyhedra with a total of n vertices and a given line ℓ, in $O(k^2 n \log n)$ time, we can compute a partitioning of ℓ into intervals so that from the viewpoints in the same interval, the silhouette arrangement remains combinatorially invariant. When the polyhedra constitute a terrain with k mountains and ℓ is a vertical line, such a subdivision for silhouette maps can be computed in $O(kn\beta_4(k)\log^2 n)$ time. The silhouette structure can be updated in output sensitive time at each event.*

In [7], a method is presented to solve the problem of answering a ray shooting query for a ray emanating from a given line. Similarly, we can build such a ray-shooting data structure once we compute the points at which a combinatorial change occurs in the silhouette map. Constructing such ray shooting structure takes roughly $O(k^2 n)$ and $O(kn)$ space and preprocessing time, for general polyhedral scenes and terrains, respectively.

6. Open questions

The main open questions are:

1. For silhouettes-union changes, the lower and upper bounds under algebraic motions are $\Omega(k^2 n + n^2)$ and $O(kn^2)$, respectively. There is a gap of $O(k)$.
2. We have described an algorithm to compute all the changes for a viewpoint moving on a given line in about $O(k^2 n)$ time. Can we do it in an on-line manner, for example, in a kinetic data structures framework?

Acknowledgements

We thank Danny Halperin for suggesting the proof of Theorem 3.1 and T.M. Murali for a helpful discussion. We thank Tamara Munzner and an anonymous SODA referee for pointing us to some relevant literature. Finally we thank the antonymous CGTA referees for many important comments and suggestions for improving the paper.

References

[1] B. Aronov, M. Sharir, The common exterior of convex polygons in the plane, Comput. Geom. Theory Appl. 8 (1997) 139–149.
[2] B. Aronov, H. Bronnimann, D. Halperin, R. Schiffenbauer, On the number of views of polyhedral scenes, in: Proc. of the Japanese Conference on Discrete and Computational Geometry (JCDCG), in: Lecture Notes in Computer Science, vol. 2098, Springer-Verlag, 2000, pp. 81–90.
[3] G. Barequet, C. Duncan, M. Goodrich, S. Kumar, M. Pop, Efficient perspective-accurate silhouette computation, in: Proc. 15th Annu. ACM Sympos. Comput. Geom., 1999, pp. 417–418.
[4] J. Basch, L.J. Guibas, J. Hershberger, Data structures for mobile data, in: Proc. 8th ACM–SIAM Sympos. Discrete Algorithms, 1997, pp. 747–756.
[5] M. de Berg, D. Halperin, M. Overmars, M. van Kreveld, Sparse arrangements and the number of views of polyhedral scenes, Internat. J. Comput. Geom. Appl. 7 (1997) 175–195.
[6] M. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf, Computational Geometry, second ed., Springer-Verlag, 2000.
[7] M. Bern, D. Dobkin, D. Eppstein, R. Grossman, Visibility with a moving point of view, Algorithmica 11 (1994) 360–378.

[8] F. Benichou, G. Elber, Output sensitive extraction of silhouettes from polygonal geometry, in: Proc. The Seventh Pacific Conference on Computer Graphics and Applications, 1999, pp. 60–69.
[9] W. Bouma, G. Vaněček Jr., Velocity-based collision detection, in: Graphics Gems V, Academic Press, 1995.
[10] R. Cole, M. Sharir, Visibility problems for polyhedral terrains, J. Symbolic Comput. 7 (1989) 11–30.
[11] A. Efrat, L.J. Guibas, O.A. Hall-Holt, L. Zhang, On incremental rendering of silhouette maps of a polyhedral scene, in: Proc. 11th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA), 2000, pp. 910–917.
[12] F.C. Crow, Shadow algorithms for computer graphics, Comput. Graph. 11 (1977) 242–248.
[13] X. Gu, S.J. Gortler, H. Hoppe, L. McMillan, B.J. Brown, A.D. Stone, Silhouette mapping, Tech. Rep. TR-1-99, Department of Computer Science, Harvard University, 1999.
[14] A. Hertzmann, D. Zorin, Illustrating smooth surfaces, in: Proc. 27th Annual Conference on Computer Graphics and Interactive Techniques, 2000, pp. 517–526.
[15] L. Kettner, E. Welzl, Contour edge analysis for polyhedron projections, in: Geometric Modeling: Theory and Practice, in: Focus in Computer Graphics, Springer-Verlag, 1997, pp. 379–394.
[16] H.P. Lenhof, M. Smid, Maintaining the visibility map of spheres while moving the viewpoint on a circle at infinity, Report MPI-I-92-102, Max-Planck-Institut Inform., 1992.
[17] D. Luebke, C. Erikson, View-dependent simplification of arbitrary polygonal environments, in: Proc. SIGGRAPH 1997, pp. 199–208.
[18] K. Mulmuley, Hidden surface removal with respect to a moving point, in: Proc. 23rd Annu. ACM Sympos. Theory Comput., 1991, pp. 512–522.
[19] M. Pop, C.A. Duncan, G. Barequet, M.T. Goodrich, W. Huang, S. Kumar, Efficient perspective-accurate silhouette computation and applications, in: Proc. Symposium Computational Geometry 2001, pp. 60–68.
[20] P.V. Sander, X. Gu, S.J. Gortler, H. Hoppe, J. Snyder, Silhouette clipping, in: Proc. SIGGRAPH 2000, pp. 327–334.
[21] M. Sharir, P.K. Agarwal, Davenport–Schinzel Sequences and Their Geometric Applications, Cambridge University Press, New York, 1995.