On the worst-case complexity of the silhouette of a polytope

Helmut Alt, Marc Glisse, Xavier Goaoc

On the worst-case complexity of the silhouette of a polytope

Helmut Alt [1] Marc Glisse [2] Xavier Goaoc [3]

May 13, 2003

Abstract

We give conditions under which the worst-case size of the silhouette of a polytope is sub-linear. We provide examples with linear size silhouette if any of these conditions is relaxed. Our bounds are the first non-trivial bounds for the worst-case complexity of silhouettes.

1 Introduction

Given a viewpoint, the apparent boundary of a polyhedron, or *silhouette*, is the set of edges incident to a visible face and an invisible one; a face whose supporting plane contains the viewpoint is considered invisible. With this definition, the silhouette is a simple closed curve on the surface of the polyhedron that separates visible and invisible faces.

Silhouettes appear in various problems in computer graphics, such as hidden surface removal [4] or shadow computations [1, 2]. The most common geometric primitives are polyhedra, so a better understanding of the size of their silhouette yields direct improvement in the theoretical complexity of algorithms in computer graphics.

Practical observations, supported by an experimental study by Kettner and Welzl [5], suggest that the number of silhouette edges of a polyhedron is usually much smaller than the total number of edges. However, only one theoretical result backs up these observations: in the same paper, Kettner and Welzl proved that a polyhedral approximation of a sphere with Hausdorff distance ϵ has $\Theta(\frac{1}{\epsilon})$ edges, and a random orthographic silhouette has size $\Theta(\frac{1}{\sqrt{\epsilon}})$.

In this paper, we investigate the worst case size of the silhouette of a polytope observed under orthographic projection. We prove that some classes of regular polytopes have orthographic silhouettes with sub-linear complexity in the worst-case. We also give examples with linear size silhouette when any of our regularity conditions is relaxed, hence showing that they are minimal. We work on polytopes to avoid the problems induced by self-occlusion, and use orthographic projection because of its relative simplicity.

Our approach is to consider the orthogonal projection of the polytope on a plane, since the boundary of the projected polygon is the projection of the silhouette. We measure the length of the boundary of the orthogonal projection of a silhouette, which we call its *apparent length*; then we show that a triangulated fat object with n edges of length $\Theta(1)$ has silhouettes with $O(\sqrt{n})$ apparent length. Then we derive bounds on the number of silhouette edges, using an additional condition on the repartition of the directions of the edges.

This paper is organized as follows. First, in Section 2, we review some examples of ill-shape polytopes with silhouettes of linear complexity. Next, Section 3 studies the apparent length of the silhouette, and Section 4 relates it to the number of silhouette edges. Last, Section 5 discusses extensions and applications of our results.

2 First examples

The goal of this paper is to find conditions under which polytopes have sub-linear silhouette in the worst-case. This section examines three examples of ill-shaped polytopes with silhouette of linear complexity, and isolates what makes them exhibit this behavior.

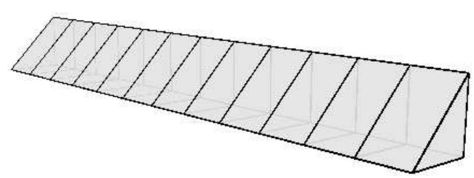

Figure 1: An elongated polytope.

[1] Email: alt@inf.fu-berlin.de

[2] LORIA - École Normale Supérieure, 45, rue d'Ulm, 75005 Paris, France. Email: marc.glisse@ens.fr

[3] LORIA - Nancy 2, 615, rue du Jardin Botanique, B.P. 101, 54602 Villers-les-Nancy, France. Email: xavier.goaoc@loria.fr

The example of Fig. 1 is characteristic of polytopes much longer along one dimension than along the others. This kind of behavior can be ruled out by considering *fat* polytopes, i.e. polytopes such that the ratio of the radius of the smallest enclosing to the largest enclosed sphere is $O(1)$.

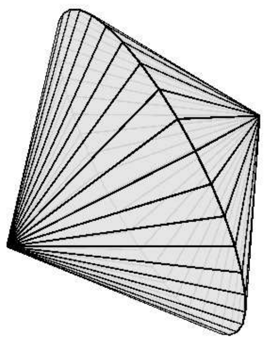

Figure 2: A polytope with uneven edges.

Our second example (see Fig. 2) illustrates the impact of the length on the edges over the silhouette. The ratio of the length of the longest edge to the length of the smallest is of order n, the total number of points. To avoid such behaviors, we require our polytopes to have *bounded edges*, i.e. that all edges are of length $\Theta(1)$.

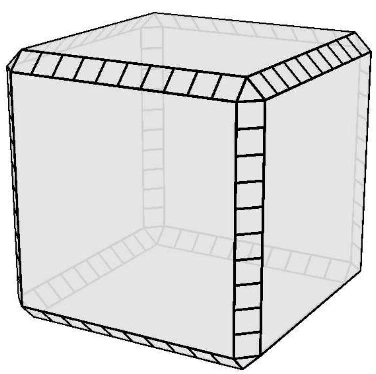

Figure 3: A polytope with a face of large complexity.

Our last example, in Fig. 3, exhibits a linear-size silhouette due to faces with order n edges. We therefore consider polytopes with faces of bounded complexity. Without loss of generality, we simply assume that our polytopes are triangulated.

This set of conditions is minimal in some sense: each of the previous examples satisfies all but one condition.

In summary, in the rest of this paper we consider *triangulated fat* polytopes with *bounded edges.*

3 Apparent length

Recall that the apparent length of the silhouette is defined as the length of the orthogonal projection of the silhouette on a plane. In this section, we relate the number of edges of the silhouette to its apparent length.

We first recall a classical result on measures of convex sets. A proof can be found in [6][1].

Lemma 1 *Let $\mathcal{O}$ and $\mathcal{O}'$ be two convex objects in $\mathbb{R}^2$ (resp. $\mathbb{R}^3$) such that $\mathcal{O}$ contains $\mathcal{O}'$. Then the length (resp. area) of $\partial\mathcal{O}$ is larger than that of $\partial\mathcal{O}'$.*

For a polytope $\mathcal{P}$, let $\mathcal{A}(\mathcal{P})$ denote its surface area, and $\mathcal{L}(\mathcal{P})$ be the maximum apparent length of its silhouettes. The following lemma relates those two quantities.

Lemma 2 *If $\mathcal{P}$ is a fat polytope, then*

$$\mathcal{L}(\mathcal{P}) = \Theta(\sqrt{\mathcal{A}(\mathcal{P})}).$$

Proof. Let r be the radius of the biggest enclosed sphere of $\mathcal{P}$, and λr be the radius of the smallest enclosing sphere. Since $\mathcal{P}$ is fat, λ is $\Theta(1)$.

First, we apply Lemma 1 to $\mathcal{P}$ and its biggest enclosed sphere, and to $\mathcal{P}$ and its smallest enclosing sphere. This yields that $\mathcal{A}(\mathcal{P}) = \Theta(r^2)$. Next, consider an orthogonal projection of $\mathcal{P}$. Each of the two spheres projects into a circle of the same radius. Since the projection of $\mathcal{P}$ is convex, we can apply Lemma 1 to these circles and the boundary of that projection, and obtain that the length of that boundary is $\Theta(r)$. Taking the maximum over all possible orthogonal projections, we obtain that $\mathcal{L}(\mathcal{P}) = \Theta(r)$. It follows that $\mathcal{L} = \Theta(\sqrt{\mathcal{A}})$. □

The next lemma bounds the area of a polytope with bounded edges.

Lemma 3 *If $\mathcal{P}$ is a triangulated polytope with bounded edges, then $\mathcal{A}(\mathcal{P}) = O(n)$.*

Proof. Since the polytope has bounded edges, the area of any of its triangles is $O(1)$. By Euler's formula, a triangulated polytope with n edges has $O(n)$ triangles, and the result follows. □

[1]In fact, the proof in [6] is much more general than our statement, and applies to any Minkowski measure, in any dimension.

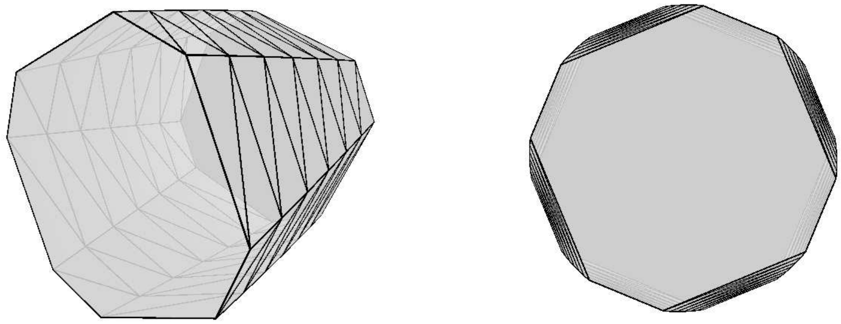

Figure 4: A triangulated fat polytope with bounded edges and a linear-size silhouette (the front and back faces were not triangulated for clarity).

We can conclude with the following corollary, directly deduced from Lemmas 2 and 3.

Corollary 4 *If $\mathcal{P}$ is a triangulated fat polytope with n bounded edges, then $\mathcal{L}(\mathcal{P}) = O(\sqrt{n})$.*

4 Complexity of the silhouette

This section uses Corollary 4 to measure the complexity of the silhouette.

To exploit the upper bound on the apparent length of the silhouette, we simply bound from below the contribution of silhouette edges to the apparent length. However, the contribution of an edge can be arbitrarily small, as it can be parallel to the direction of projection, and a triangulated fat polytope with bounded edges can have a linear number of such silhouette edges, as shown in Fig. 4. Thus, we need to bound from above the number of silhouette edges that can be close to the direction of projection.

We give two distinct additional conditions that ensure a sub-linear size for the silhouette. The first one is a local condition.

Theorem 5 *Let ϵ be some positive real number and $\mathcal{P}$ be a triangulated fat polytope with n bounded edges such that any two incident edges make an angle in the interval $[\epsilon, \pi - \epsilon]$. Then, any silhouette of $\mathcal{P}$ has $O(\sqrt{n})$ edges.*

Proof. Let us choose a viewing direction $\vec{\delta}$. Since any two incident edges make an angle in the interval $[\epsilon, \pi - \epsilon]$, two consecutive silhouette edges contribute $\Omega(\epsilon)$ to the apparent length of the silhouette. Since $\mathcal{L}(\mathcal{P}) = O(\sqrt{n})$, by Corollary 4, it follows that the number of silhouette edges is $O(\sqrt{n})$. Note that the constant in the O depends on ϵ. $\square$

The second condition is global and corresponds to a regular repartition of the directions of the edges of the polytope.

Theorem 6 *Let $\mathcal{P}$ be a triangulated fat polytope with n bounded edges such that for any direction $\vec{\delta}$, the number of edges of $\mathcal{P}$ making an angle smaller than $\Theta(n^{-1/6})$ with $\vec{\delta}$ is $O(n^{2/3})$. Then any silhouette of $\mathcal{P}$ has $O(n^{2/3})$ edges.*

Proof. Let us fix a direction $\vec{\delta}$, and let α be a real number. We count separately the silhouette edges that make an angle greater than α with $\vec{\delta}$, and the others, and find the value of α yielding the best trade-off.

If we represent the set of directions by a unit sphere, the directions that make an angle smaller than α with $\vec{\delta}$ form a spherical cap of area $\Theta(\alpha^2)$. Given that the sphere can be covered by $\Theta(1/\alpha^2)$ such spherical caps, and that the directions of the n edges are distributed over the sphere, one of the caps has to contain $\Omega(\alpha^2 n)$ edge directions. It means that, for some viewing direction, there are $\Omega(\alpha^2 n)$ edges that make an angle less than α. Thus, the best we can ask is that the number of silhouette edges having a negligible contribution to the apparent length is $O(\alpha^2 n)$.

Let k denote the number of silhouette edges that make an angle greater than α with $\vec{\delta}$. The contribution of these k edges to the apparent length of the silhouette is $\Omega(k\alpha)$. Thus, $k = O(\mathcal{L}/\alpha)$, and applying Corollary 4, we get that $k = O(\sqrt{n}/\alpha)$.

If we ask that at most $O(\alpha^2 n)$ edges of the polytope make an angle less than α with any given direction, then the complexity of the silhouette is bounded from above by

$$O\left(\frac{\sqrt{n}}{\alpha} + \alpha^2 n\right).$$

The best trade-off we can achieve is to choose

$$\frac{\sqrt{n}}{\alpha} = \Theta(\alpha^2 n),$$

which means $\alpha = \Theta(n^{-1/6})$. In that case, the number of silhouette edges is $O(n^{2/3})$, and the regular distribution assumption is the one mentioned in the statement of the theorem. □

Note that the proof of Theorem 6 establishes a more general result: a weaker condition on the repartition of the directions of the edges still yields a sub-linear bound on the complexity of the silhouette, which is in between $O(n^{2/3})$ and $O(n)$.

5 Discussion

This section discusses our results, giving extensions as well as possible applications.

To begin with, notice that, for all the results of Sections 3 and 4, the fatness assumption can be weakened. In fact, Lemma 2 holds for any polytope $\mathcal{P}$ with bounded edges that satisfies

$$d(\mathcal{P})^2 = O(\mathcal{A}(\mathcal{P}))$$

where $d(\mathcal{P})$ is its diameter. This is equivalent to having a fat orthogonal projection with the same diameter, or, intuitively, to be fat along at least two dimensions.

Next, to extend our approach to the perspective case, one has to deal with two distinct issues. First, the distance from the object to the viewpoint has to be taken into account. When the viewpoint is far from the polytope, the perspective case should behave as the orthographic case. But when the viewpoint is close to the polytope, the perspective projection introduces a lot of distortion: the length of the projection of a silhouette edge greatly depends on its distance to the center of the view, so the apparent length and the number of silhouette edges are no longer directly related. Second, since the apparent length of an edge does not depend only on its direction, the global hypothesis on the distribution of the directions of the edges has to be adjusted accordingly.

The results of this paper are only a first step toward the understanding of the complexity of silhouettes, but they still have promising applications.

A first application is the computation of shadow boundaries. Drettakis and Duguet [1, 2] propose a solution based on a *visibility skeleton* restricted to the visual events generated by a punctual light source. In their detailed report [2], they show that their algorithm has complexity $O(ns_n)$, where n is the size of the polyhedron that casts a shadow, and s_n the size of its silhouette. Even the orthographic case is of interest, since it corresponds to a light source at infinity, a simple sun model for instance.

A second application is hidden surface removal, which has a long history as a problem difficult to address practically [3]. A solution proposed by Efrat et al. [4] is to render separately the silhouettes of the objects, and the single-object regions. They estimate the number of combinatorial changes to the rendered silhouettes of polytopes when the viewpoint moves along a line or an algebraic curve. Depending on the motion, this number depends either linearly or quadratically on the silhouette complexity, which they bound from above by the complexity of the polytope. Extension of our work to the perspective case would thus yield a direct improvement of their bounds.

Acknowledgements

The authors wish to thank Olivier Devillers for mentioning this problem during the Workshop on 3D Global Visibility in 2000 at Bonifacio, Corsica, and Sylvain Petitjean for his helpful comments. Part of the research was done during the First McGill-INRIA Workshop on Geometry in Graphics at McGill Bellairs Research Institute, in 2002.

References

[1] F. Duguet and G. Drettakis. Robust Epsilon Visibility. SIGGRAPH 2002, pp 567-575.

[2] F. Duguet. Implémentation robuste du squelette de visibilité. Masters Thesis (in french), 2001.

[3] S. E. Dorward. A survey of object-space hidden surface removal Internat. J. Comput. Geom. Appl., volume 4, pp. 325–362, 1994.

[4] A. Efrat, L. J. Guibas, O.A. Hall-Holt, and L. Zhang. On incremental rendering of silhouette maps of a polyhedral scene. SODA 2000, pp 910-917.

[5] L. Kettner and E. Welzl. Contour edge analysis for polyhedron projections. *Geometric Modeling: Theory and Practice*, Springer, pp. 379-394, 1997.

[6] L. Santalò. Integral probability and geometric probability. in *Encyclopedia of mathematics and its applications*, volume 1. Addison-Wesley, Reading, MA, 1979.