ELSEVIER

Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT®

Computational Geometry ●●● (●●●●) ●●●–●●●

Computational Geometry
Theory and Applications

www.elsevier.com/locate/comgeo

On the number of views of translates of a cube and related problems ☆

Boris Aronov [a,*], Robert Schiffenbauer [a], Micha Sharir [b,c]

[a] *Department of Computer and Information Science, Polytechnic University, Brooklyn, NY 11201, USA*
[b] *School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel*
[c] *Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA*

Received 23 October 2003; accepted 31 October 2003

Communicated by P. Agarwal

Abstract

It is known that a general polyhedral scene of complexity n has at most $O(n^6)$ combinatorially different orthographic views and at most $O(n^9)$ combinatorially different perspective views, and that these bounds are tight in the worst case. In this paper we show that, for the special case of scenes consisting of a collection of n translates of a cube, these bounds improve to $O(n^{4+\varepsilon})$ and $O(n^{6+\varepsilon})$, for any $\varepsilon > 0$, respectively. In addition, we present constructions inducing $\Omega(n^4)$ combinatorially different orthographic views and $\Omega(n^6)$ combinatorially different perspective views, thus showing that these bounds are nearly tight in the worst case. Finally, we show how to extend the upper and lower bounds to several classes of related scenes.

Keywords: Visibility; Aspect graphs; Orthographic views; Perspective views; Fat objects; Polyhedral terrains; Arrangements; Envelopes; Combinatorial geometry

☆ Work on this paper by Boris Aronov and Micha Sharir has been supported by a joint grant from the US–Israeli Binational Science Foundation. Work by Boris Aronov has also been supported by NSF Grants CCR-99-72568 and ITR CCR-00-81964. Work by Micha Sharir has also been supported by a grant from the Israel Science Fund (for a Center of Excellence in Geometric Computing), by NSF Grants CCR-97-32101 and CCR-00-98246, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

* Corresponding author.

E-mail addresses: aronov@cis.poly.edu (B. Aronov), rschiff@cis.poly.edu (R. Schiffenbauer), michas@post.tau.ac.il (M. Sharir).

0925-7721/$ – see front matter © 2003 Published by Elsevier B.V.
doi:10.1016/j.comgeo.2003.10.001

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 cgt708 P.2 (1-14)
by:JOL p. 2

1. Introduction

Aspect graphs [14], which arise in visibility theory, are data structures that incorporate information about all the possible views of an object or collection of objects in a given scene; here, we consider *opaque polyhedral* objects only. Aspect graphs are useful, for example, in object classification, in which the identity of an unknown object is established by comparing an available subset of its views with the views of a series of known objects and ascertaining the closest match. Roughly speaking, a *view* is the line drawing resulting from the projection of object features visible from a given viewpoint onto a two-dimensional viewing plane or pair of planes. In the literature, there are two distinct models under which views may be generated. Under *orthographic projection* each viewpoint lies on the sphere at infinity and all lines of sight emanate from the viewpoint in the same direction. The viewpoint is defined by two parameters: θ, its *longitude*, or angle of rotation about the vertical axis, and φ, its *azimuth*, or angle from the positive vertical axis. The view is the projection of the visible portions of object edges and visible object vertices onto a plane orthogonal to the lines of sight. Under *perspective projection*, each viewpoint lies in free space in $\mathbf{R}^3$ and lines of sight emanate from the viewpoint in all directions. The viewpoint is defined by its x-, y-, and z-coordinates. The view is the projection of the visible portions of object edges and visible object vertices onto some pair of parallel planes containing the viewpoint in the slab between them.

One variation on this theme is that sometimes the view is defined to contain only ‘significant’ object features; for example, in some cases only those visible (portions of) edges and vertices belonging to the silhouette of an object with respect to a given viewpoint are projected to form the view from that viewpoint [7].

In either model, viewpoint space is partitioned into maximal connected regions such that the views from the viewpoints in any region are *isomorphic*. That is, the views, when considered as labeled, embedded, undirected planar graphs, are all topologically equivalent [14]. Under perspective projection the (three-dimensional) maximal connected regions are separated by planar or quadric surfaces. Under orthographic projection the (two-dimensional) maximal connected regions are separated by geodesic or quadratic curves which are the intersections of the planar or quadric surfaces under perspective projection with the sphere at infinity. The curves or surfaces separating regions of viewpoints with topologically equivalent views are referred to as *critical* curves or surfaces. Each consists of those viewpoints for which there exists some *critical event* occurring in the associated view such that the views corresponding to viewpoints immediately on one side of the curve or surface are non-isomorphic to the views corresponding to viewpoints immediately on the other side. It can be shown [9] that critical events are of two types only. *EV events* occur due to the alignment along some line of sight of an object vertex and a point on an object edge (both of which are visible) so that the projections of the vertex and edge intersect at a point in the view. *EEE events* occur due to the alignment of three visible points on three object edges along some line of sight so that the projections of these edges intersect at a point in the view. Clearly, for a general polyhedral scene with n features (vertices, edges and faces), EV events induce at most $O(n^2)$ critical curves or surfaces, while EEE events induce at most $O(n^3)$ critical curves or surfaces. For our purposes we consider an EV event to be a special type of EEE event, involving the alignment of one endpoint of each of two edges (adjacent to the vertex) and a point on a third edge.

We say that a critical event is *occluded* at a viewpoint when it is rendered invisible from that viewpoint due to the imposition of an object face (along the line of sight at which the event would have occurred) between the viewpoint and at least one of the edges inducing the event. Viewpoints at which a critical event is occluded are not part of the critical curves or surfaces induced by that event.

Given a critical event, an *event occlusion endpoint* (*EOE point*) [8] is a viewpoint such that, for any $\varepsilon > 0$, a ball with center at that viewpoint and radius ε will contain both viewpoints from which the event is occluded and viewpoints from which the event is not occluded. This implies that at any EOE point there exists a line of sight along which four scene edges align; the three edges which induce the associated critical event and a fourth edge adjacent to the object face causing the occlusion.

In the orthographic case, critical curves terminate either abruptly at EOE points or naturally because the edges inducing their associated critical events are of finite length. It can be shown that, in the worst case, the number of EOE points dominates the number of points at which the curves terminate naturally. In the perspective case, critical surfaces are either bounded by EOE points or are bounded naturally, again, because the edges inducing the critical event are of finite length. It can be shown that, in the worst case, the number of critical surface edges (vertices) formed by EOE points dominates the number of critical surface edges (vertices) at which the surfaces terminate naturally.

The arrangement [10] of critical curves or surfaces in viewpoint space induced by any polyhedral scene is called the *viewpoint space partition*, a structure dual to the aspect graph [13]. It follows that, in the orthographic case, a bound on the number of vertices in the viewpoint space partition can be found by bounding the number of EOE points plus the number of points at which the relative interiors of two critical curves intersect. Further, in the perspective case, a bound on the number of vertices can be found by bounding the number of points at which a critical surface edge formed by EOE points intersects the relative interior of a second critical surface plus the number of points at which the relative interiors of three critical surfaces intersect (we note that the number of all other vertices, those adjacent to critical surface edges at which the surfaces terminate naturally, is $O(n^3)$, and that this is dominated by the bound we shall prove for the perspective case). In either case a bound on the complexity of the viewpoint space partition is obtained. This, in turn, provides a bound on the total number of non-isomorphic views induced by the scene. For a general polyhedral scene of complexity n, Plantinga and Dyer [13] have shown this, in the worst case, to be $\Theta(n^6)$ under orthographic projection and $\Theta(n^9)$ under perspective projection.

In this paper we shall be mostly interested in scenes consisting of (bounded) convex, *fat* polyhedra. A (bounded) convex polyhedron is *fat* [11] if the ratio of the radius of the largest ball contained within the polyhedron to the radius of the smallest ball containing the polyhedron is bounded away from zero. Intuitively, such objects possess no arbitrarily long, skinny parts.

The objects that populate our scenes include (translates of) cubes, *rectilinear near-unit-cubes*, *skyscraper terrains*, *zonohedra* and arbitrary convex centrally symmetric polyhedra. A cube is a fat object. We define a *near-unit-cube* to be a parallelepiped whose edge lengths lie in the interval from one up to a constant $m \geqslant 1$. A near-unit-cube cannot therefore be *too* long and skinny or *too* flat. A *rectilinear* polyhedron (or polyhedral surface) is such that each of its edges is parallel to one of the coordinate axes. A *skyscraper terrain*, which will be defined more precisely in Section 4, is, essentially, a connected infinite rectilinear polyhedral surface with features that can be long and skinny in the vertical direction only. A *zonohedron* is a convex polyhedron formed by taking the Minkowski sum of finitely many line segments [6].

Our results. In the main result of this paper, we establish that for scenes consisting of a collection of n pairwise disjoint translates of a cube the maximum possible number of non-isomorphic views is in fact lower than the bounds given above. Alternatively, the worst-case complexity of the viewpoint space

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) cgt708 P.4 (1-14)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 by:JOL p. 4

partition induced by such scenes is lower than for general polyhedral scenes. In other words, the *effects of occlusion* become significant for these restricted scenes.

Thus we present the first known non-trivial bounds on the number of views of a scene *consisting exclusively of fat objects*. Little is known regarding bounds on the number of views induced by more general scenes of fat objects. We emphasize that our results are for a particularly simple scene of this type, and that, in the general case, the problem remains open.

Agarwal and Sharir [2] and de Berg et al. [4] have previously demonstrated the existence of additional restricted classes of polyhedral scenes of complexity n (see below) for which the bounds on the number of views are lower than in the general case.

Related work. A great deal of research has focused on visibility questions in general and combinatorial and algorithmic issues related to aspect graphs in particular. Plantinga and Dyer [13] offered constructions showing that the trivial upper bounds of $O(n^6)$ and $O(n^9)$ for the complexity of the viewpoint space partition induced by general polyhedral scenes of complexity n under orthographic and perspective projection (respectively) are in fact tight in the worst case. Snoeyink [16] showed that the bound under orthographic projection continues to be tight in the case of scenes consisting solely of rectilinear (long and skinny) parallelepipeds. De Berg et al. [4] improved the bounds of Plantinga and Dyer to $O(n^4k^2)$ under orthographic projection and to $O(n^6k^3)$ under perspective projection in the special case of a scene consisting of k pairwise disjoint convex polyhedra with total complexity n. Recently, Aronov et al. [3] provided a lower bound construction which establishes that these bounds are also tight in the worst case. De Berg et al. [4] also improved the upper bound of Plantinga and Dyer to $O(n^5 \cdot 2^{c(\log n)1/2})$ (for a constant $c > 0$) under orthographic projection in the case of a general polyhedral terrain. In addition, they demonstrated a lower bound of $\Omega(n^5\alpha(n))$ (where $\alpha(n)$ is the slowly growing inverse Ackermann function), thus showing that the upper bound is nearly tight. Agarwal and Sharir [2] improved the upper bound of Plantinga and Dyer to $O(n^{8+\varepsilon})$ (where $\varepsilon > 0$ may be selected as small as desired by an appropriate choice of the implied constant) under perspective projection in the case of a general polyhedral terrain. De Berg et al. [4] demonstrated a lower bound of $\Omega(n^8\alpha(n))$, thus showing that the upper bound is nearly tight. See [14] for a more complete survey of recent research efforts related to aspect graphs.

Outline of the paper. In Section 2 we demonstrate upper bounds of $O(n^{4+\varepsilon})$ under orthographic projection and $O(n^{6+\varepsilon})$ under perspective projection, for any $\varepsilon > 0$, for the complexity of the viewpoint space partition induced by scenes consisting of n pairwise disjoint translates of a cube. Thus the maximum possible number of views associated with such scenes is significantly lower than in the general case. In Section 3 we present constructions for which the number of views is $\Omega(n^4)$ under orthographic projection and $\Omega(n^6)$ under perspective projection, thus nearly closing the gap between the upper and lower bounds. We note here that these bounds show that, in the worst case, a relatively large viewpoint space partition complexity is already inherent even in very simple scenes of fat objects. In Section 4 we show how to extend the upper bound results to the union of possibly overlapping rectilinear near-unit-cubes and to pairwise disjoint translates of a zonohedron. We also show that the upper bounds hold for a skyscraper terrain and indicate constructions similar in principle to those exhibited in Section 3 for the lower bounds under orthographic and perspective projection. Finally, we note that the upper bound results also apply to arbitrary convex centrally symmetric polyhedra when only silhouette views are considered.

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) cgt708 P.5 (1-14)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 by:JOL p. 5

2. Upper bounds

Let C be a collection of n pairwise disjoint translates of a fixed cube P. We write $C = \{P_i = a_i \oplus P\}_{i=1,\ldots,n}$, (where '$\oplus$' denotes the Minkowski sum with the singleton $\{a_i\}$) and refer to the vector a_i as the *translation vector* of P_i, for $i = 1, \ldots, n$. We wish to bound the complexity of the viewpoint space partition induced by C by counting certain classes of its vertices, as described above.

2.1. Orthographic views

Consider the case of orthographic views. Let S denote the unit sphere of directions. For each $\mathbf{u} \in S$, consider the orthographic projection $C(\mathbf{u})$ of the cubes in C onto some plane orthogonal to $\mathbf{u}$. The family $C(\mathbf{u})$ consists of n translates of the projection $P(\mathbf{u})$ of P. Specifically, $C(\mathbf{u}) = \{a_i(\mathbf{u}) \oplus P(\mathbf{u})\}_{i=1,\ldots,n}$, where $a_i(\mathbf{u})$ is the projection of a_i (again, '$\oplus$' denotes the Minkowski sum, this time in the plane).

We need to bound the number of orientations $\mathbf{u}$ at which one of the following two types of events occurs:

(i) There exist a quadruple of indices i_1, i_2, i_3, i_4 and a ray λ in direction $-\mathbf{u}$, such that λ touches an edge of each of the four cubes $P_{i1}, P_{i2}, P_{i3}, P_{i4}$ (in the order $P_{i4}, P_{i3}, P_{i2}, P_{i1}$, with λ emanating from a point on the edge of P_{i4}), and λ does not intersect the interior of any cube. The number of orientations at which this type of event occurs yields the number of EOE points in the viewpoint space partition.

(ii) There exist two distinct triples of indices (possibly with common elements) i_1, i_2, i_3 and j_1, j_2, j_3, and two distinct rays λ, λ' in direction $-\mathbf{u}$, such that (a) λ touches an edge of each of the three cubes P_{i1}, P_{i2}, P_{i3} (in the order P_{i3}, P_{i2}, P_{i1}, with λ emanating from a point on the edge of P_{i3}), (b) λ' touches an edge of each of the three cubes P_{j1}, P_{j2}, P_{j3} (in the order P_{j3}, P_{j2}, P_{j1}, with λ' emanating from a point on the edge of P_{j3}), and (c) neither λ nor λ' intersects the interior of any cube. The number of orientations at which this type of event occurs yields the number of intersection points between the relative interiors of two critical curves in the viewpoint space partition.

The projection $P_i(\mathbf{u})$ of any translate P_i of P, for a direction $\mathbf{u} \in S$, has a silhouette which is generally a convex centrally symmetric hexagon. Three additional edges of P_i are visible, and appear as internal edges within $P_i(\mathbf{u})$. Each of them is a translate of two edges of the silhouette of $P_i(\mathbf{u})$, and together they partition $P_i(\mathbf{u})$ into three parallelograms. Three additional edges of P_i are invisible when viewed in direction $\mathbf{u}$.

Note that, in both types of events, only the edge(s) containing the endpoint(s) of the appropriate ray(s) (the edge of P_{i4} in a type (i) event, or the edges of P_{i3} and of P_{j3} in a type (ii) event) can be interior in the respective projection(s); all other edges must be silhouette edges.

Fix one translate $P_0(\mathbf{u}) = a_0(\mathbf{u}) \oplus P(\mathbf{u})$, and fix an edge (either a silhouette edge or an internal edge) $e_0 = e_0(\mathbf{u})$ of $P_0(\mathbf{u})$. For any other translate $P'(\mathbf{u}) = a'(\mathbf{u}) \oplus P(\mathbf{u})$, consider the intersection $I' = e_0(\mathbf{u}) \cap P'(\mathbf{u})$. As is easily verified, I' is an interval along e_0 which *contains an endpoint of* e_0. This follows from the observation that the length of any cross section of a convex centrally symmetric polygon in a direction parallel to a side of it, is always at least as large as the length of that side, and that $e_0(\mathbf{u})$ is, or has the same length as and is parallel to, a side of the silhouette.

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) cgt708 P.6 (1-14)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 by:JOL p. 6

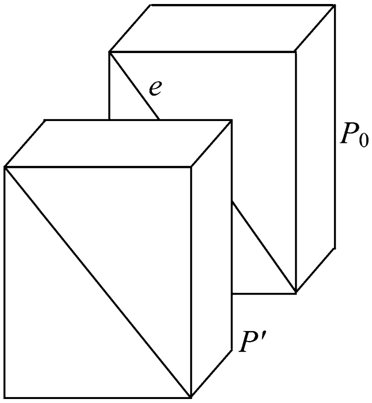

Fig. 1. The edge e of P_0 is partially hidden by a nearer translate P' but no endpoint of e is hidden.

Remark. The argument just given remains valid as long as $e_0(\mathbf{u})$ is (a translate of) a silhouette edge. This holds for cubes, as noted above, and, more generally, for zonohedra, but may fail for general (centrally symmetric) convex polyhedra; see Fig. 1 for an example.

Denote the endpoints of $e_0(\mathbf{u})$ by $p_0(\mathbf{u})$ and $q_0(\mathbf{u})$. Parametrize (the line containing) e_0 by orienting it from p_0 to q_0, and by representing a point x on it by its signed distance from p_0 (so q_0 has a positive representation).

For each translate P' as above, define two (partially-defined) real-valued functions, $F^{(e_0)}_{P'}$, $G^{(e_0)}_{P'}$ on S, so that $F^{(e_0)}_{P'}(\mathbf{u})$ is $\max I'$ (in the parametrization of e_0), provided that (a) $p_0 \in I'$ and (b) P' is in front of P_0 as viewed from $\mathbf{u}$; otherwise, $F^{(e_0)}_{P'}(\mathbf{u})$ is undefined. Similarly, $G^{(e_0)}_{P'}(\mathbf{u})$ is $\min I'$, provided that $q_0 \in I'$ and P' is in front of P_0 as viewed from $\mathbf{u}$; otherwise it is undefined.

It is an easy exercise to verify that, with an appropriate parametrization of S, the functions $F^{(e_0)}_{P'}$ and $G^{(e_0)}_{P'}$ are of *constant description complexity* [15], i.e., the graph of each function is a semialgebraic set defined by a Boolean combination of a constant number of equations and inequalities involving polynomials of constant maximum degree.

Let $E^-_{e_0}$ denote the upper envelope of the functions $F^{(e_0)}_{P'}$, and let $E^+_{e_0}$ denote the lower envelope of the functions $G^{(e_0)}_{P'}$.

Let $\mathbf{u}$ be an orientation of type (i), with a corresponding quadruple P_{i1}, P_{i2}, P_{i3}, P_{i4} of translates of P, and respective contact edges e_1, e_2, e_3, e_4. Then, by definition and construction, $\mathbf{u}$ is the projection on S either of a vertex of $E^-_{e_4}$, or of a vertex of $E^+_{e_4}$, or of an intersection of an edge of one envelope with a facet of the other. In other words, the event corresponds to a vertex of the *sandwich region* S_{e4} enclosed between the two envelopes $E^-_{e_4}$ and $E^+_{e_4}$. Since each of these is the envelope of $n-1$ bivariate functions of constant description complexity, the results of Agarwal et al. [1] (see also [12]) imply that the complexity of such a sandwich region, and thus also the number of type (i) events that involve e_4 as their furthest contact edge (with respect to $\mathbf{u}$), is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Repeating this argument for each edge of every translate of P, we conclude that the number of events of type (i) is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.

Consider next the analysis of events of type (ii). For each edge e of any translate of P, denote the projection of (the edges and vertices of) the sandwich region S_e onto S by M_e. Then, again by definition

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) cgt708 P.7 (1-14)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 by:JOL p. 7

and construction, an event of type (ii) occurs at $\mathbf{u}$ if there are two cube edges e_1, e_2, so that an edge of M_{e1} crosses an edge of M_{e2} at $\mathbf{u}$. (It is possible that $e_1 = e_2 = e$, in which case the crossing is between the projections of an edge of E_e^- and of an edge of E_e^+.)

For any fixed pair of edges e_1, e_2, the complexity of the overlay of M_{e1} and of M_{e2} is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This is a consequence of the following result, which extends the analysis of overlays given in [1,12].

Lemma 2.1.1. *Let F_1, F_2, G_1, G_2 be four collections of bivariate functions of constant description complexity each of size at most n. Let S_F denote the sandwich region between the upper envelope of F_1 and the lower envelope of F_2, and let S_G denote the sandwich region between the upper envelope of G_1 and the lower envelope of G_2. Let M_F (respectively, M_G) denote the projection of (the edges and vertices of) S_F (respectively, S_G) onto the xy-plane. Then the complexity of the overlay of M_F and of M_G is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Proof. Let Q denote the overlay of the maximization diagram [10] of F_1, the maximization diagram of G_1, the minimization diagram [10] of F_2, and the minimization diagram of G_2. By the results of [1,12], the complexity of Q is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Let q be a crossing point of an edge e of M_F and of an edge e' of M_G. By definition, e is either an edge of the maximization diagram of F_1, or an edge of the minimization diagram of F_2, or the projection of an edge of intersection between the upper envelope of F_1 and the lower envelope of F_2. Similarly, e' is of one of three corresponding types, defined in terms of G_1 and G_2.

If e is of one of the first two types and so is e' then q is a vertex of Q. On the other hand, suppose that both e and e' are of the third type, where e (respectively, e') is the projection of a portion of an intersection curve between the graphs of some $f_1 \in F_1$ and $f_2 \in F_2$ (respectively, of some $g_1 \in G_1$ and $g_2 \in G_2$). Let τ be the cell of Q that contains q. By construction, τ is fully contained in a single cell of each of the four maximization or minimization diagrams of the respective F_1, F_2, G_1, G_2. This is easily seen to imply that τ uniquely determines the four functions f_1, f_2, g_1, g_2 that define q, which in turn implies that τ can contain only $O(1)$ crossing points q of the above kind. Similar reasoning applies when e is of the third kind and e' is of one of the two former kinds, or vice versa. Since the number of cells τ is $O(n^{2+\varepsilon})$, the lemma follows. □

Multiplying the bound provided by Lemma 2.1.1 by the number $O(n^2)$ of pairs of edges e_1, e_2, we obtain an overall bound of $O(n^{4+\varepsilon})$ for the number of type (ii) events. We thus obtain the main result of this section:

Theorem 2.1.1. *The number of combinatorially different orthographic views of a collection of n pairwise disjoint translates of a cube in $\mathbf{R}^3$ is $O(n^{4+\varepsilon})$, for any $\varepsilon > 0$.*

We will see below (Section 3) that this bound is nearly tight in the worst case.

2.2. Perspective views

Consider next the case of perspective views. For each $\mathbf{z} \in \mathbf{R}^3$, consider the central projection $C(\mathbf{z})$ of the cubes in C from $\mathbf{z}$ onto some pair of parallel planes containing $\mathbf{z}$ in the slab between them.

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24
cgt708
P.8 (1-14)
by:JOL p. 8

Without loss of generality, assume that the planes are parallel to a facet of P, and that this facet is horizontal. Technically, we prefer this projection over the more natural projection onto a sphere centered at $\mathbf{z}$, because the images in our projection are convex polygons. It suffices to analyze the changes that occur in just one of these planes. Some cubes may project onto both planes, and then both projections are unbounded polygons. As $\mathbf{z}$ passes through a horizontal plane that contains a facet of some translate P_i, the projection of P_i on one plane starts or stops being nonempty. In what follows we omit the analysis of the effect of these changes on the number of views, since they do not affect the asymptotic bound that we are going to derive.

The collection $C(\mathbf{z})$ (on the fixed projection plane) consists of n projections of P, each of which is a (possibly unbounded) convex polygon. Similar to the orthographic case, each projection contains some additional visible projected edges in its interior.

We need to bound the number of points $\mathbf{z}$ at which one of the following two types of events occurs:

(i) There exist a quadruple of indices i_1, i_2, i_3, i_4 and a triple of indices j_1, j_2, j_3 (they may share elements, but the triple is not fully contained in the quadruple), and two distinct segments s, s' having $\mathbf{z}$ as a common endpoint, such that (a) s touches an edge of each of the four cubes $P_{i1}, P_{i2}, P_{i3}, P_{i4}$ (in that order), with the other endpoint of s lying on the edge of P_{i4}, and s does not intersect the interior of any cube; and (b) s' touches an edge of each of the three cubes P_{j1}, P_{j2}, P_{j3} (in that order), with the other endpoint of s' lying on the edge of P_{j3}, and s' does not intersect the interior of any cube. The number of points at which this type of event occurs yields the number of intersection points between a critical surface edge formed by EOE points and the relative interior of a second critical surface in the viewpoint space partition.

(ii) There exist three distinct triples of indices (possibly with common elements) $i_1, i_2, i_3, j_1, j_2, j_3$, and k_1, k_2, k_3, and three distinct segments s, s', s'' with $\mathbf{z}$ as a common endpoint, such that s satisfies the property in (i) with respect to the triple i_1, i_2, i_3 (with its other endpoint lying on an edge of P_{i3}), and s', s'' satisfy this property with the triples j_1, j_2, j_3, and k_1, k_2, k_3, respectively. The number of points at which this type of event occurs yields the number of intersection points among the relative interiors of three critical surfaces in the viewpoint space partition.

Fix one of the projections $P_0(\mathbf{z})$, and fix an edge $e_0 = e_0(\mathbf{z})$ of $P_0(\mathbf{z})$. For any other translate $P'(\mathbf{z})$ such that P' is in front of P_0 as viewed from $\mathbf{z}$, consider the intersection $I' = e_0(\mathbf{z}) \cap P'(\mathbf{z})$. We claim that if I' is nonempty then it must be an interval along e_0 which *contains an endpoint of* e_0.

To prove the claim, let Δ_0 denote the triangle spanned by $\mathbf{z}$ and the edge f_0 of P_0 which projects to e_0. Denote the endpoints of f_0 by a and b. The claim is equivalent to asserting that if P' intersects Δ_0 then it intersects at least one of the edges za or zb of Δ_0.

Let p be a point on P' intersecting Δ_0; see Fig. 2. Let l be the line through p parallel to f_0. Note that the segment $s = P' \cap l$ is parallel to an edge f' of P' and thus has length equal to that of f'. Hence the intersection of P' and Δ_0 contains a (contiguous) portion of a segment lying on P' that is parallel to f_0 and this segment is of length (at least) as long as that of f_0. Thus P' must intersect either za or zb. This implies our claim.

Remark. The argument just given remains valid as long as $e_0(\mathbf{z})$ is the projection of (a translate of) a silhouette edge. As in the preceding section, this holds for cubes and, more generally for zonohedra, but may fail for general (centrally symmetric) convex polyhedra.

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) cgt708 P.9 (1-14)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 by:JOL p. 9

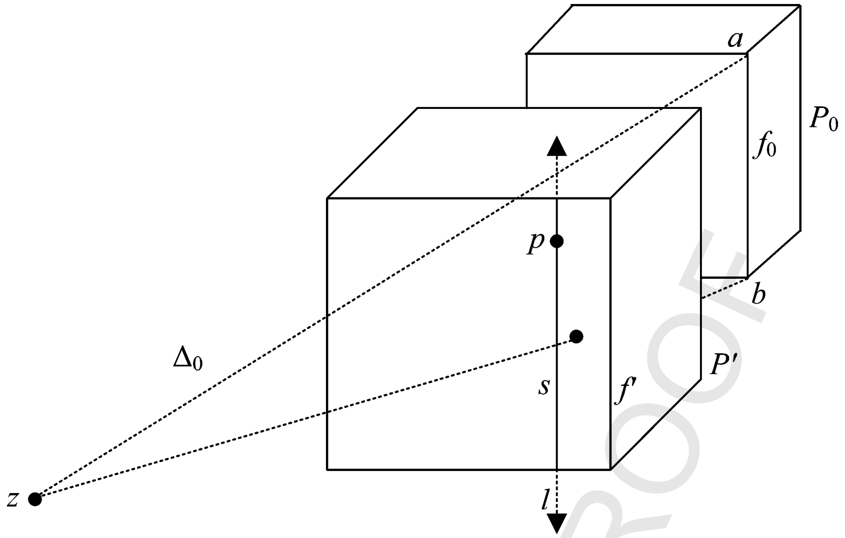

Fig. 2. A nearer translate P' that intersects Δ_0 must intersect an edge of Δ_0.

Denote the endpoints of $e_0(\mathbf{z})$ by $p_0(\mathbf{z})$ and $q_0(\mathbf{z})$, and parametrize (the line containing) e_0 as in the preceding section. For each translate P' as above, define two (partially-defined) real-valued functions, $F_{P'}^{(e_0)}$, $G_{P'}^{(e_0)}$ on $\mathbf{R}^3$, so that $F_{P'}^{(e_0)}(\mathbf{z})$ is $\max I'$ (in the parametrization of e_0), provided that (a) $p_0 \in I'$ and (b) P' is in front of P_0 as viewed from $\mathbf{z}$; otherwise, $F_{P'}^{(e_0)}(\mathbf{z})$ is undefined. Similarly, $G_{P'}^{(e_0)}(\mathbf{z})$ is $\min I'$, provided that $q_0 \in I'$ and P' is in front of P_0 as viewed from $\mathbf{z}$; otherwise it is undefined.

As in the preceding section, it is an easy exercise to verify that the functions $F_{P'}^{(e_0)}$ and $G_{P'}^{(e_0)}$ are of constant description complexity.

Let $\mathbf{z}$ be a point in 3-space of type (i), with a corresponding quadruple $P_{i1}, P_{i2}, P_{i3}, P_{i4}$ of translates of P, respective contact edges e_1, e_2, e_3, e_4, another triple P_{j1}, P_{j2}, P_{j3} of cubes, and their respective contact edges e'_1, e'_2, e'_3. Then, by definition and construction, $\mathbf{z}$ is an intersection point between the projection on $\mathbf{R}^3$ of an edge of the sandwich region S_{e4} and the projection of a 2-face of the sandwich region $S_{e'3}$.

Since these are sandwich regions between envelopes of $n-1$ trivariate functions of constant description complexity, the recent results of Koltun and Sharir [12] can be used to deduce that the number of such intersection points is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.

Indeed, this bound is an immediate consequence of the following extension of Lemma 2.1.1 to the case of trivariate functions:

Lemma 2.2.1. *Let $F_1, F_2, G_1, G_2, H_1, H_2$ be six collections of trivariate functions of constant description complexity each of size at most n. Let S_F denote the sandwich region between the upper envelope of F_1 and the lower envelope of F_2, and define S_G, S_H analogously for the collections G_1, G_2 and H_1, H_2, respectively. Let M_F (respectively, M_G, M_H) denote the projection of (the faces, edges and vertices of) S_F (respectively, S_G, S_H) onto the xyz-hyperplane. Then the complexity of the overlay of M_F, M_G and M_H is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 cgt708 P.10 (1-14)
by:JOL p. 10

Proof. An easy generalization to the case of trivariate functions of the proof of Lemma 2.1.1, using the near-cubic bound on the complexity of the overlay of any constant number of minimization (or maximization) diagrams of trivariate functions of constant description complexity, as given in [12]. □

Remark. Clearly, the same bound holds if we consider only two sandwich regions, rather than three.

Applying Lemma 2.2.1, and multiplying the resulting bound by the number $O(n^2)$ of pairs of furthest (from $\mathbf{z}$) contact edges e_4, e_3', we obtain that the number of type (i) points is $O(n^{5+\varepsilon})$, for any $\varepsilon > 0$.

Analysis of points of type (ii) is also straightforward, and proceeds in much the same way as above. It applies Lemma 2.2.1 to the three sandwich regions, each arising for the furthest edge of contact in each of the triples of translates of P that are involved in the event. We omit the further easy details. This yields a bound of $O(n^{3+\varepsilon})$ for the number of points of type (ii) associated with a fixed triple of furthest contact edges, and, multiplying by the number $O(n^3)$ of such triples of edges, we obtain an overall bound of $O(n^{6+\varepsilon})$, for any $\varepsilon > 0$. We thus obtain the main result of this section:

Theorem 2.2.1. *The number of combinatorially different perspective views of a collection of n pairwise disjoint translates of a cube in* $\mathbf{R}^3$, *is* $O(n^{6+\varepsilon})$, *for any* $\varepsilon > 0$.

We will see below (Section 3) that this bound is nearly tight in the worst case.

3. Lower bounds

We now present lower bound constructions inducing $\Omega(n^4)$ and $\Omega(n^6)$ different views under orthographic and perspective projection, respectively, as follows.

For orthographic projection, let R be the set of viewpoints in a small rectangular region just below and to the right of the origin of the plane $y = +\infty$. Place a collection of $\Theta(n)$ (rectilinear) cubes along the negative y-axis so that they appear, from R, to be lined up one behind the other and so that the upper edge of each cube face for which the outward normal vector is in the $+y$ direction is just visible above the upper edge of the corresponding face of the cube immediately in front of it (*group* (a), Fig. 3).

Next, place a collection of $\Theta(n)$ (rectilinear) cubes along the positive y-axis so that they appear, from R, to be lined up one behind the other and so that the top right vertex of each cube face for which the outward normal vector is in the $+y$ direction appears to be slightly below and to the right of the top right vertex of the corresponding face of the cube immediately in front of it (*group* (b), Fig. 3).

Each line of sight emanating from any viewpoint in R and passing through the top right vertex (as specified above) of a cube in *group* (b) will be tangent to the cube at that vertex.

The edges and vertices specified above combine to create $\Theta(n^2)$ EV events each of which induces a critical surface intersecting the plane $y = +\infty$ along a horizontal line in R. The cubes may be positioned so that these critical curves are pairwise disjoint in R and so that the distance between neighboring curves is arbitrarily smaller than the lengths of the curves themselves.

Finally, copy and translate the cubes in *group* (b) to form *group* (d), and copy, translate and rotate the cubes in *group* (a) to form *group* (c) (Fig. 3). This induces a second set of $\Theta(n^2)$ critical curves in R orthogonal to the first set. The cubes may be positioned so that the curves in the second set intersect each of the curves in the first set. This forms a grid of $\Theta(n^2)$ by $\Theta(n^2)$ critical curves in the region R on

ARTICLE IN PRESS
S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24
cgt708
P.11 (1-14)
by:JOL p. 11

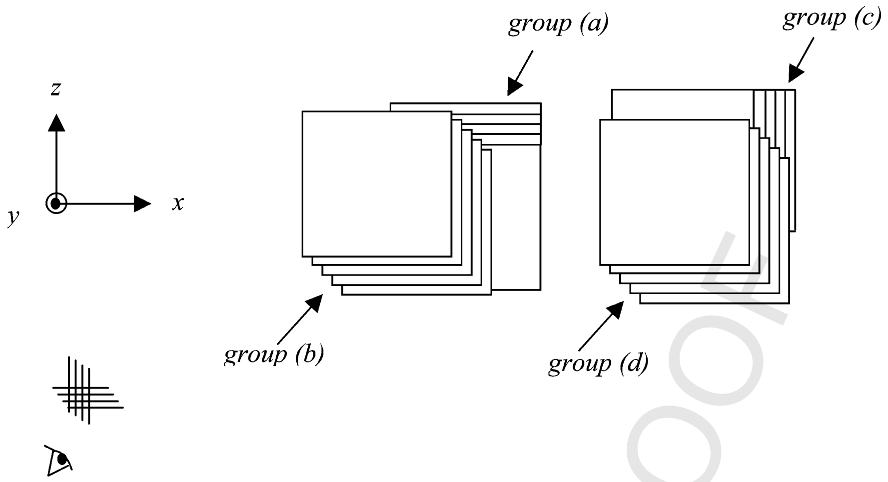

Fig. 3. The construction under orthographic projection.

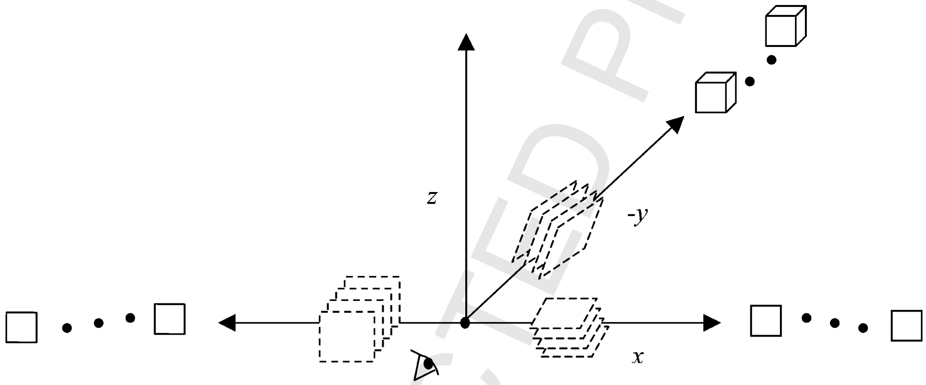

Fig. 4. The construction under perspective projection (and its induced critical surfaces).

the plane $y = +\infty$. Hence the complexity of the viewpoint space partition under orthographic projection induced by this scene is $\Omega(n^4)$.

For the perspective case, we position a copy of *groups* (a) and (b) on the positive x-axis, a copy of *groups* (c) and (d) on the negative y-axis and a *mirror image* copy of *groups* (c) and (d) (reflected through the yz-plane) on the negative x-axis. Each copy is placed sufficiently far from the origin, and is oriented so that the critical surfaces induced form a $\Theta(n^2)$ by $\Theta(n^2)$ by $\Theta(n^2)$ grid within a small parallelepiped near the origin (see Fig. 4). Thus the complexity of the viewpoint space partition induced by this scene is $\Omega(n^6)$.

4. Discussion

Pairwise disjoint near-unit-cubes. The collection of pairwise disjoint translates of a cube in the proofs of Section 2 may be replaced more generally with a collection of pairwise disjoint rectilinear *near-unit-*

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 cgt708 P.12 (1-14) by:JOL p. 12

cubes of different sizes, that is, axis-parallel parallelepipeds whose edge lengths lie in the interval from one up to a constant $m \geqslant 1$. Assume that an edge e in this new scene has length l_e for $1 \leqslant l_e \leqslant m$. We may subdivide e into m subintervals of length l_e/m ($\leqslant 1$), identifying the projection of each subinterval in turn with e_0 (in the discussions of Sections 2.1 and 2.2), and apply the analysis given in those sections to each of the m subintervals created. In particular, the assertion continues to hold that I' (as defined in Sections 2.1 and 2.2) must contain an endpoint of e_0 (we say that I' *covers* that endpoint). Thus the upper bounds presented in Section 2 remain valid for these more general scenes.

Overlapping translates of a cube. The union of a collection of n possibly overlapping translates of a cube with unit length edges, which may, in addition to convex edges, contain arbitrarily short concave edges, itself has complexity $O(n)$ [5]. Any such concave edge will be parallel to some (collection of) silhouette edges in the scene. In addition, if general position is assumed, no edge has length greater than one. Note that if a concave edge is involved in an EV or EEE event, or in the creation of an EOE point, it must be the *furthest* edge from the viewpoint along the line of sight associated with that event or EOE point.

It can be seen that the analysis of Section 2 may be applied to the union of overlapping translates of a cube in general position. In particular, it continues to hold that I' covers an endpoint of e_0, even when e_0 is the projection of a concave edge. Thus the upper bounds are applicable to these scenes also.

Overlapping near-unit-cubes. The union of a collection of n possibly overlapping rectilinear near-unit-cubes of different sizes has complexity $O(n)$. This becomes evident by observing that each near-unit-cube may be approximated as closely as desired by a constant number of translates of a cube with unit length edges (slightly perturbed so as to be in general position) and that there exist at least as many features in the new scene as there were in the original, after which the proof presented in [5] may be applied directly. Again the upper bounds are extendible to these more general scenes.

Skyscraper terrains. We consider a collection of n possibly overlapping rectilinear parallelepipeds of varying heights (*skyscrapers*), whose bases lie on the xy-plane, having the property that all edges parallel to the x- and y-axes possess lengths lying in the interval from one up to a constant $m \geqslant 1$. Edges parallel to the z-axis may be of arbitrary length. We take the boundary of the union of these parallelepipeds along with the entire xy-plane, excluding those portions of the xy-plane containing the bases. The resulting connected infinite two-dimensional polyhedral surface will be referred to as a *skyscraper terrain*.

It is not difficult to see that the analysis of Section 2 may be applied to skyscraper terrains. In particular, by virtue of the properties of a terrain (see Fig. 5), and with the usual assumption that viewpoints are restricted to points *above* the terrain only, it continues to hold that I' covers an endpoint of e_0, even when e_0 is the projection of a vertical edge. Moreover, a simple counting argument on the number of vertices in a skyscraper terrain can be used to show that its complexity is $O(n)$. Therefore, the upper bounds are extendible to skyscraper terrains (note that, for vertical edges, the analysis of Section 2 is somewhat simplified since only facets of envelopes, rather than sandwich regions, need be considered).

We note that our upper bounds of $O(n^{4+\varepsilon})$ for any $\varepsilon > 0$ under orthographic projection and $O(n^{6+\varepsilon})$ for any $\varepsilon > 0$ under perspective projection for the specific case of a skyscraper terrain represent improvements over the respective upper bounds of $O(n^5 \cdot 2^{c(\log n)^{1/2}})$ (for a constant $c > 0$) derived by de Berg et al. [4] and $O(n^{8+\varepsilon})$ derived by Agarwal and Sharir [2] for general polyhedral terrains.

We further point out that a simple modification to the first construction exhibited in Section 3 allows us to deduce that a lower bound under orthographic projection for the case of skyscraper terrains is $\Omega(n^4)$ (the modification is that we change all cubes to parallelepipeds with bases on the xy-plane).

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●) ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24 cgt708 P.13 (1-14) by:JOL p. 13

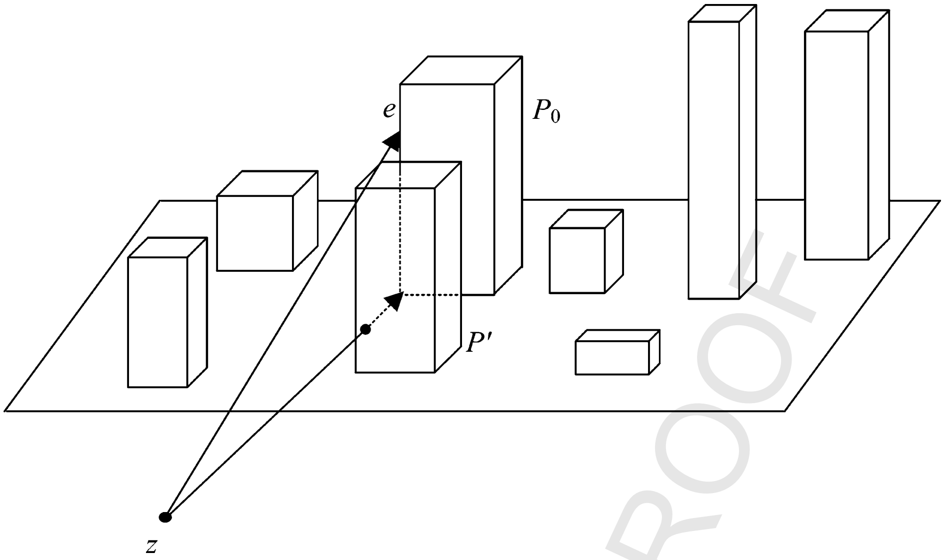

Fig. 5. In a skyscraper terrain, a nearer parallelepiped P' that partially hides a vertical edge e of P_0 must also hide the lower endpoint of e.

A similar modification to the second construction of Section 3 allows us to deduce that a lower bound under perspective projection for skyscraper terrains is $\Omega(n^6)$.

Zonohedra. As previously noted, the analysis of Section 2 holds for pairwise disjoint translates of a zonohedron with O(1) facets. This is so because the arguments given there remain valid as long as e_0 is the projection of (a translate of) a silhouette edge, which, for zonohedra, will always be the case. In particular, it continues to hold that I' covers an endpoint of e_0, when e_0 is the projection of *any* edge in the scene.

Centrally symmetric polyhedra. We also note that the analysis of Section 2 holds for pairwise disjoint translates of arbitrary convex centrally symmetric polyhedra with O(1) facets, provided that we only consider views of their silhouettes. That is, we assume that edges of the polyhedra that become internal in the projected view are not observable in the view, so critical events only involve silhouette edges. In particular, it continues to hold that I' covers an endpoint of e_0, whenever e_0 is the projection of a silhouette edge.

Fat objects. We reiterate that improving the trivial upper bounds of $O(n^6)$ and $O(n^9)$ on the complexity of the viewpoint space partition induced by *general* scenes comprised of n fat objects each of complexity O(1) remains an open problem. Other simple scenes of this type, for which there are no known non-trivial upper bounds, include, for example, disjoint translates of a simplex, disjoint translated *and scaled* copies of a cube, or disjoint translated *and rotated* copies of a cube.

References

[1] P.K. Agarwal, O. Schwarzkopf, M. Sharir, The overlay of lower envelopes in 3-space and its applications, Discrete Comput. Geom. 15 (1996) 1–13.

[2] P.K. Agarwal, M. Sharir, On the number of views of polyhedral terrains, Discrete Comput. Geom. 12 (1994) 177–182.

S0925-7721(03)00076-2/FLA AID:708 Vol.●●●(●●●)
ELSGMLTM(COMGEO):m2 v 1.177 Prn:2/12/2003; 15:24
cgt708
P.14 (1-14)
by:JOL p. 14

[3] B. Aronov, H. Brönnimann, D. Halperin, R.D. Schiffenbauer, The number of views of polyhedral scenes, in: Proceedings of the Japanese Conference on Discrete and Computational Geometry (JCDCG 2000), in: Lecture Notes In Computer Science, Vol. 2098, Tokai University, Japan, 2000, pp. 81–90.
[4] 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 (3) (1997) 175–195.
[5] J.D. Boissonnat, M. Sharir, B. Tagansky, M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, Discrete Comput. Geom. 19 (1998) 485–519.
[6] H.S.M. Coxeter, Regular Polytopes, Dover, 1973.
[7] A. Efrat, L. Guibas, O. Hall-Holt, L. Zhang, On incremental rendering of silhouette maps of a polyhedral scene, in: Proceedings of the ACM-SIAM Symposium On Discrete Algorithms, 2000, pp. 910–917.
[8] Z. Gigus, J. Canny, R. Seidel, Efficiently computing and representing aspect graphs of polyhedral objects, IEEE Trans. Pattern Anal. Machine Intell. 13 (6) (1991) 542–551.
[9] Z. Gigus, J. Malik, Computing the aspect graph for line drawings of polyhedral objects, IEEE Trans. Pattern Anal. Machine Intell. 12 (2) (1990) 113–122.
[10] D. Halperin, Arrangements, in: Handbook of Discrete and Computational Geometry, CRC Press LLC, Boca Raton, FL, 1997, pp. 389–412.
[11] M.J. Katz, K.R. Varadarajan, A tight bound on the number of geometric permutations of convex fat objects in $\mathbf{R}^d$, Discrete Comput. Geom. 26 (2001) 543–548.
[12] V. Koltun, M. Sharir, The partition technique for the overlay of envelopes, SIAM J. Comput. 32 (2003) 841–863.
[13] H. Plantinga, C.R. Dyer, Visibility, occlusion, and the aspect graph, Internat. J. Comput. Vision 5 (2) (1990) 137–160.
[14] R.D. Schiffenbauer, A survey of aspect graphs, Polytechnic University Department of Computer and Information Science Technical Rpt. TR-CIS-2001-01, available at http://cis.poly.edu/tr/tr-cis-2001-01.htm, 2001.
[15] M. Sharir, P.K. Agarwal, Davenport—Schinzel Sequences and Their Geometric Applications, Cambridge University Press, New York, 1995.
[16] J. Snoeyink, The number of views of axis-parallel objects, Algorithms Rev. 2 (1991) 27–32.