Pseudo-triangulations : Theory and Applications

Michel Pocchiola *   Gert Vegter †

1 Introduction

Pseudotriangles and pseudo-triangulations have played a key role in the recent design of two optimal visibility graph algorithms; see [1, 2]. The purpose of this paper is (1) to give three new applications of these concepts to 2-dimensional visibility problems, and (2) to study realizability questions suggested by the pseudotriangle-pseudoline duality; see Figure 1.

Our first application is related to the ray-shooting problem in the plane: preprocess a set of objects into a data structure such that the first object hit by a query ray can be computed efficiently. In section 3 we show that for a scene of n objects, where the objects are pairwise disjoint convex sets with m 'simple' arcs in total, one can obtain $O(\log m)$ query time using $O(m+k)$ storage, where $k = O(n^2)$ is the size of the visibility graph of the set of obstacles. Previous solutions use $\Theta(n^2)$ storage for a similar query time. (We refer to [3, 4] for a bibliography on the ray-shooting problem.) An other feature of our data structure is that it can be used to compute in $O(h \log m)$ time the h objects visible from a query point in a query interval of directions; we mention also that our technique can be extended to scenes of non-convex obstacles. Due to lack of space these latter two points will only be developed in the full version of the paper. Our proof is based on the 'right-to-left' property of the so-called greedy pseudo-triangulation (see section 2); it uses dynamic point-location data structures for plane graphs and persistent data structures. (We refer to [4] for a bibliography related to these data structures.) The two other applications are related to optimal covering problems. In section 4 we show that (1) computing a lighting set with worst case minimal size (i.e., $4n-7$, as shown in [5]) for a set of n disjoint convex sets reduces in $O(n)$ time to computing a pseudo-triangulation; (Using the Koebe Representation Theorem we also give a practical characterization of all cases in which $4n-7$ lighting points are required.) and (2) computing a polygonal cover with worst case minimal size (i.e., with no more than $6n-9$ sides and $3n-6$ slopes, as shown in [6]) for a set of n disjoint convex sets, reduces in $O(n)$ time to computing a pseudo-triangulation. Our polygonal cover algorithm is simpler than the algorithm[1] described by M. de Berg [3]. In section 5 we examine realizability questions suggested by the pseudotriangle-pseudoline duality. One of the main questions concerning arrangements of pseudolines is realizability by arrangements of straight lines (also called stretchability): given a configuration of pseudolines, is it isomorphic to an arrangement of straight lines? It is known that 'most' arrangements of pseudolines are not stretchable, and that the realizability question is NP-hard. (See [7, 8] for background material and recent developments on this topic.) A set of pseudotriangles whose dual image is isomorphic to a given arrangement of pseudolines will be called a *realization* of this arrangement. We show that any arrangement of pseudolines can be realized by a set of pseudotriangles. It is tempting to conjecture that it can even be realized by a

*Département de Mathématiques et Informatique, Ecole normale supérieure, ura 1327 du Cnrs, 45 rue d'Ulm 75230 Paris Cedex 05, France (pocchiola@dmi.ens.fr)

†Dept. of Math. and Comp. Sc, University of Groningen P.O.Box 800, 9700 AV Groningen, The Netherlands (gert@cs.rug.nl)

[1] This latter algorithm, attributed to R. Wenger, applies only to (convex) polygons and achieves an $O(n)$ size for the cover but not worst case size optimality. Finally we mention that polygonal covers with few vertices can be used to answer efficiently depth order queries on terrains; see M. de Berg [3, pages 132–133].

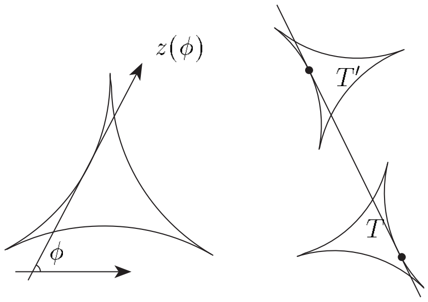

Figure 1: A pseudotriangle is a simply connected subset T of the plane, such that (i) its boundary ∂T consists of three smooth convex curves that are tangent at their endpoints, (ii) T is contained in the triangle formed by the three endpoints of these convex curves. For each ϕ, $0 \le \phi \le 2\pi$, the boundary of a pseudotriangle has exactly one directed tangent line that makes an angle of ϕ with the positive x-axis. The curve of directed tangent lines to T is described by its ϕ-parametrization $z(\phi) : \mathcal{R} \to \mathcal{S}^2$, where $z(\phi) = -z(\phi + \pi)$, and $z(\phi)$ is a line with slope ϕ. (We identify the point (a, b, c), with $c \ne \pm 1$, on the 2-sphere $\mathcal{S}^2 = \{(x, y, z) \in \mathcal{R}^3 \mid x^2 + y^2 + z^2 = 1\}$ with the directed line with equation $ax + by + c = 0$ and direction $u = (-b/r, a/r) \in \mathcal{S}^1$, $r = (a^2 + b^2)^{1/2}$.) Therefore the dual curve $\phi \mapsto \{z(\phi), -z(\phi)\}$ of the pseudotriangle is a pseudoline in the projective plane $\mathcal{P}^2 = \mathcal{S}^2/\{x, -x\}$, quotient of the 2-sphere by its antipodal isomorphism. According to [1], two disjoint pseudotriangles share exactly one common tangent line; in other terms the set of dual curves of a set of pairwise disjoint pseudotriangles is an arrangement of pseudolines, called the dual arrangement of the pseudotriangles.

set of *disjoint* pseudotriangles, but so far we have only been able to prove this for a large class, of size 2^{cn^2}, for some poistive constant c, of arrangements of n pseudolines. A byproduct of our study is a lower bound for the number of visibility graphs/complexes of configurations of n convex obstacles of the form 2^{cn^2}, for some positive constant c. If we only consider convex objects of *degree* $d = O(n^\alpha)$, for some fixed $0 \le \alpha < 1$, (i.e., whose boundaries consist of at most d arcs of complexity $O(1)$) the lower bound is of the form $2^{\Omega(dn \log n)}$.

2 Background material

Let $\mathcal{O} = \{O_i\}$ be a finite set of n pairwise disjoint bounded closed convex subsets of the Euclidean plane (obstacles for short). We assume that the boundary of O_i is a ϕ-curve[2] given by its ϕ-parametrization $z_i(\phi)$ as the product of m_i 'simple' ϕ-arcs, in the sense that the ≤ 2 common tangent lines of two ϕ-arcs are computable in constant time. We set $m = \sum_{i=1}^n m_i$. A *maximal* (*minimal*) point of an obstacle is a boundary point at which the tangent line is horizontal, such that the obstacle lies below (above) this tangent line. An *extremal* point is either a maximal or a minimal point. A *bitangent* is a closed line segment whose supporting line is tangent to two obstacles at its endpoints. It is called free if it lies in free space, i.e., the complement of the union of the interior of the obstacles. We denote by B the set of free bitangents, and by k its cardinality. Bitangents in B are oriented upward.

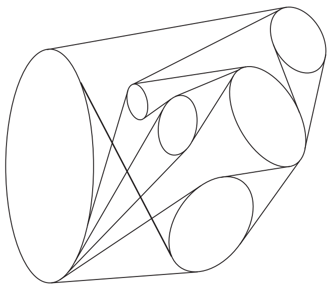

Figure 2: The right-to-left property.

A *pseudo-triangulation* G is a maximal subset of non-crossing free bitangents in B. The boundaries of the obstacles and the bitangents in G induce a regular cell decomposition of the plane, still called a pseudo-triangulation and denoted by $\mathcal{H}(G)$. According to [1], the bounded free faces of $\mathcal{H}(G)$ are pseudotriangles, their number is $2n - 2$, and the cardinality of G is $3n - 3$. The *greedy* pseudo-triangulation $G_0 = \{b_1, b_2, \ldots, b_{3n-3}\}$ is defined as follows : (1) b_1 has minimal slope in B; (2) b_{i+1} has minimal slope in the subset of bitangents in B disjoint from $b_1, b_2, \ldots, b_i$. According to [2], the pseudo-triangulation $\mathcal{H}_0 = \mathcal{H}(G_0)$ is computable in $O(m + n \log m)$ time, and verifies the remarkable 'right-to-left' property.

Theorem 1 (Right-to-left property) *[2] For all $b \in G_0$, and all $t \in B$ crossing b, the slope of t is greater than the slope of b, i.e., t pierces b from its right side to its left side.* □

The endpoints of the bitangents in B subdivide the boundaries of the obstacles into a set of arcs; these

[2]A smooth closed curve in the Euclidean plane is called a ϕ-curve, if for each ϕ, $0 \le \phi \le 2\pi$, the curve has exactly one point where its tangent makes an angle of ϕ with the positive x-axis. Such a curve is described by its ϕ-parametrization $z(\phi) : \mathcal{R} \to \mathcal{R}^2$, where $z(\phi) = z(\phi + 2\pi)$ and the unit tangent vector is $u(\phi) = (\cos\phi, \sin\phi)$ for all $\phi \in \mathcal{R}$. By definition $z([\phi_1, \phi_2])$, with $0 < \phi_2 - \phi_1 < \pi$, is called a ϕ-arc.

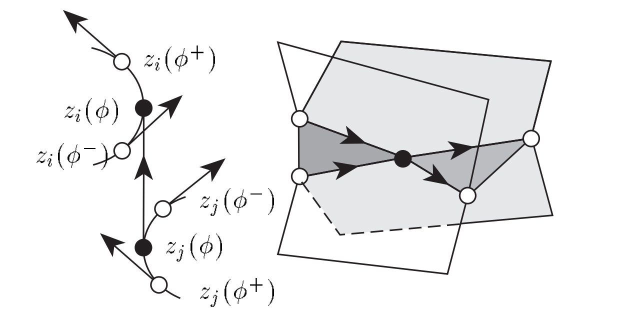

Figure 3: Neighborhood of the visibility complex at a bitangent with direction $u(\phi)$; ϕ^+ (ϕ^-) refers to $\phi + \epsilon$ ($\phi - \epsilon$) for an infinitesimal $\epsilon \in \mathcal{R}^+$.

arcs and the bitangents are the edges of the *visibility graph* of the set of obstacles. The pseudotriangulation $\mathcal{H}_0$ induces a 'partition' of B that will be useful in the sequel: For each pseudotriangle T let T^* be the slope-increasing sequence of bitangents in B whose initial or terminal point lies on ∂T, and let T^*_{Ini} be the subsequence of bitangents in T^* with initial point on ∂T; finally for $b \in G$ incident along its right side to T, we denote by $\Omega(b)$ the sequence of bitangents of T^*_{Ini} that cross b. We assume that each of these sequences of bitangents is represented by a doubly-linked list.

The structure of a visibility graph is better described in 'dual space' via the notion of *visibility complex*. A *ray* is a pair $(p, u) \in \mathcal{R}^2 \times \mathcal{S}^1$. The point p in the plane is called the origin of the ray, and the unit vector u is called its direction. For a point p in the plane we are interested in the point-obstacle (i.e., a point on $\cup\mathcal{O}$) that we can see from p in a certain direction u in $\mathcal{S}^1$. This point is called the *forward point-view* along the ray (p, u) (the *backward point-view* along the ray (p, u) is the forward point-view of the *opposite* ray, $(p, -u)$). The obstacle containing the point-view is called the *view* (from p). We denote by $\gamma(+O_i)$ ($\gamma(-O_i)$) the curve of rays $(z_i(\phi), u(\phi))$ ($(z_i(\phi), -u(\phi))$) emanating from and tangent to O_i, oriented along increasing values of ϕ.

The visibility complex X is a cell-decomposition of the quotient space V of the set of rays by the equivalence relation $\sim$, defined by $(p, u) \sim (q, u)$ if (p, u) and (q, u) have the same forward point-view. (A point in V is still called a ray.) Its 0-cells (=vertices) are the intersection points of (the images under the canonical map $\mathcal{R}^2 \times \mathcal{S}^1 \to V$ of) the curves in $\gamma(\pm\mathcal{O})$ (therefore we have a 2-1 correspondence between the vertices of X and the bitangents in B), its 1-skeleton is supported by the curves in $\gamma(\pm\mathcal{O})$ (therefore we have a 2-1 correspondence between the 1-skeleton of X and the set of edges of the visibility graph), its 2-skeleton is supported by the set of rays with origins on the obstacles' boundaries, and its 3-cells are the sets of rays with origins in the obstacles' interiors. Modulo the addition of two obstacles at infinity, the poset of cells of X ordered by the inclusion relation of their closures is an abstract polytope of rank 4; the vertex-figure of a vertex is the face poset of a 3-dimensional simplex. (See Figure 3.)

We will represent the visibility complex X by the set of planar subcomplexes[3] $X(O_i)$, whose underlying spaces $V(O_i)$ is the space of rays with backward view O_i. (Since O_i is convex, $X(O_i)$ is planar.) Each $X(O_i)$ is augmented with a point location data structure so that given a ray in $V(O_i)$ its forward view can be computed in $O(\log m)$ time. The whole representation uses $O(k)$ space.

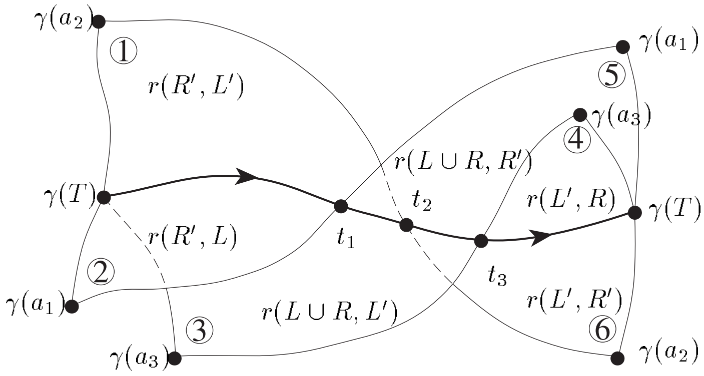

Figure 4: The visibility complex (restricted to upward rays) $X(T)$ of a red pseudotriangle T with cusp points a_i consists of 6 planar patches that correspond to sets of rays emanating from and ending on specific chains of T. Let $r(C, C')$ be the set of rays emanating from chain C and ending on chain C' then patch 1 is $r(R', L')$, etc. (The symbols a_i, R, R', L, L' refer to Figure 5. $\gamma(T)$ is the curve of upward rays emanating and tangent to T, and $\gamma(a_i)$ is the curve of upward rays with origin a_i.)

The definition of the visibility complex extends in a natural way to the case of non-convex obstacles. However cusp points and inflection points give rise to new types of vertices in the visibility complex. In the sequel we use the set of visibility complexes $X(T)$ of the pseudotriangles T in $\mathcal{H}_0$. We refer to Figures 6 and 4 for a description of T and $X(T)$. (Here the obstacle is the exterior of the pseudotriangle).

[3] A subset W of V is said to be planar if the canonical map $W \to \mathcal{S}^2$ which associates with the ray (p, u) the directed line through p and direction u, is one-to-one.

3 The ray shooting problem

Theorem 2 *A set $\mathcal{O}$ of n pairwise disjoint convex obstacles with m simple arcs in total can be stored in a data structure of size $O(k)$ such that the forward view along a query ray can be computed in time $O(\log m)$; here $k = O(n^2)$ is the size of the visibility graph of the set of obstacles.*

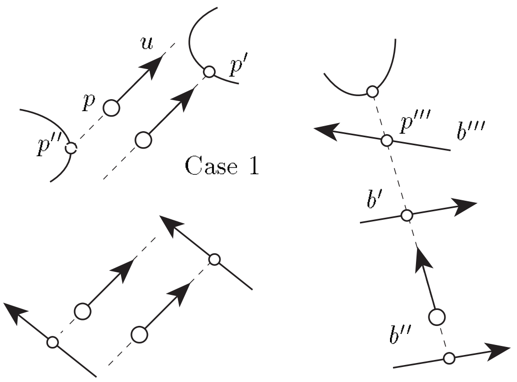

Figure 5: The three cases of the ray-shooting algorithm.

The idea behind the construction of the ray shooting data structure is based on the observation that if the backward view is known then the ray shooting problem reduces to a point location problem in the planar subcomplex associated with the known backward view. But how can we compute the backward view .... without computing the forward view? The idea is to add line segment obstacles that play the role of obstacles only along the backward directions. The right-to-left property of the greedy pseudo-triangulation G_0 is all that we need to make this idea working.

For C a convex chain of $\mathcal{H}_0$ (i.e., an alternating sequence of arcs and bitangents of $\mathcal{H}_0$ without cusp points) let $V(C)$ to be the (2-dimensional) planar space of upward rays emanating from C and pointing toward free space. The forward view mapping with respect to the obstacles *and the left sides of the bitangents in* G_0 induces a partition of $V(C)$, denoted by $X(C)$. For each bitangent $b \in G_0$ we will define a convex chain[4] $C(b)$ of $\mathcal{H}_0$ that contains the left side of b. Our ray-shooting data structure consists of the set of planar maps $X(\mathcal{O})$ and the set of planar maps $X(C(G_0))$, each augmented with a point location data structure. We explain now how a ray shooting query reduces in $O(\log m)$ time to $O(1)$ point location queries in $O(1)$ of these maps.

The algorithm proceeds as follows. Let $r = (p, u)$ be the query ray directed, w.l.o.g, upward, i.e., $u \in \mathcal{S}^1_+$. We start by computing the forward and backward point-views along r, denoted by p' and p'' respectively, in the pseudo-triangulation $\mathcal{H}_0$, i.e., p' (p'') is the first point-obstacle (i.e., in $\cup\mathcal{O}$) or point-bitangent (i.e., in $\cup G_0$) that is visible from p along the direction u ($-u$). This can be done in $O(\log m)$ time after a suitable preprocessing of the visibility complexes of the pseudotriangles in $\mathcal{H}_0$. We distinguish several cases. **Case 1.** p' (or p'') is a point-obstacle, say on obstacle O; in that case the problem reduces to locating a point in the planar map $X(O)$. **Case 2.** p' and p'' ly on bitangents b' and b'', respectively. We subdivide this case into two subcases. **Case 2.1.** p' (or p'') lies on the left (right) side of b' (b''), with respect to p. In that case the problem reduces to locating a point in $X(O)$, for some obstacle O that depends only on b' or b''; see Lemma 1. **Case 2.2.** p' and p'' lies on the right and left sides of b' and b'', respectively. In that case we compute the forward point-view p''' along the ray (p', u) in $X(b')$. If p''' lies on an obstacle we are done, if it lies on a bitangent, say b''', (necessarily on its left side) we restart the algorithm with the ray (p''', u); the crucial point is that the backward view of this ray in $\mathcal{H}_0$ is now the right side of b''', and consequently we are in case 1 or 2.1 of our algorithm.

Of course we have to show that the whole set of planar maps $X(C(b))$ with $b \in G_0$, each map being augmented with a planar point-location data structure, can be represented by a data structure with $O(k)$ size. To this end, we are going to define a partial order $<$ on G_0, and, for each down-set[5] A of $(G_0, <)$, (1) a 2-dimensional cell complex X_A, (2) a partition of X_A in planar subcomplexes $\{X_A(C_j)\}$ whose underlying spaces are spaces of rays that emanate from convex chains $\{C_j\}$, called the canonical chains of A, such that for all $b \in \max_< A$ there is a unique chain C_j ($= C(b)$) that contains the left side of b. Then given an unrefinable chain of down-sets of $(G_0, <)$: $A_0 = G_0 \supset A_1 \cdots \supset A_{3n-3} = \emptyset$ we will show that $X_{A_{i+1}}$ (and its partition) can be computed from X_{A_i} (and its partition) in $O(k_i)$ time; here k_i is the cardinality of $\Omega(b_i)$, where $b_i = A_i \setminus A_{i+1}$. In this way we can store the whole collection of planar maps $X_{A_i}(C_j)$, each augmented with a point location data structure, in $O(k)$ space using a persistent data structure. Before defining this partial order we justify the reduction claimed in case 2.1 of our algorithm.

[4] We use a chain $C(b)$ instead of b itself to achieve an $O(k + m)$ size for our ray-shooting data structure.

[5] A down-set A of $(G_0, <)$ is a subset of G_0 such that if $b \in A$ and $b' < b$ then $b' \in A$.

Ray shooting inside a pseudotriangle. A pseudotriangle is said to be *red* (*green*) if its extremal point is a maximal (minimal) point. Let T be green pseudotriangle with extremal point m and cusp points a_1, a_2 and a_3. Walking in clockwise order around its boundary, starting from its minimal point m, we find successively the convex chains ma_1, a_1a_2, a_2a_3, and a_3m, respectively denoted by R, L', R' and L and called its canonical chains, as illustrated in Figure 6. Similar notations are used for red pseudotriangles.

Lemma 1 *For a red (green) pseudotriangle the chain R (L) is free of bitangents.*
Let T be a red pseudotriangle of $\mathcal{H}_0$. The forward (backward) view function is a constant function on patches $r(R' \cup L, L'), r(L', R)$ ($r(R', L), r(R, R' \cup L')$).
Let T be a green pseudotriangle of $\mathcal{H}_0$. The backward (forward) view function is a constant function on patches $r(R', L' \cup R), r(L, R')$ ($r(R, L'), r(R' \cup L', L)$).

Proof. Assume that there is a bitangent b on chain R of a red pseudotriangle T. Let T' be the pseudotriangle incident along the left side of b. The common tangent line of T and T' is the supporting line of a free bitangent that crosses b from left-to-right, a contradiction with the right-to-left property. A similar argument applies to the chain L of a green pseudotriangle.
Let T be a red pseudotriangle and consider the patch $r(R' \cup L, L')$. We prove the result by contradiction. Assume that the forward view function along rays in $r(R' \cup L, L')$ is not constant. In that case there is an upwardly directed line segment with initial point a on the chain $R' \cup L$, that pierces L' from left to right, and terminal point b on some obstacle O_i; the segment is tangent at b to O_i. Let M_i be the maximal point of O_i. Consider the curve C with initial point the first cusp point a_1 of T, that runs along R' or L towards the point a, then along the line segment $[a, b]$, and then along the boundary of O_i from b to M_i, its terminal point. The shortest path from a_1 to M_i homotopy equivalent to C contains necessarily a bitangent emanating from R' or L and piercing L' from left-to-right. This is a contradiction with the right-to-left property of the greedy pseudo-triangulation. A similar argument applies to the other cases. □

Acyclic orientation of the greedy pseudo-triangulation. By definition, the canonical orientation of a pseudotriangle in $\mathcal{H}_0$ is given by the following rules concerning the orientation of its canonical chains R, R', L and L' : (1) R' and L' are oriented upward; (2) R is oriented upward or downward depending on whether T is green or red; (3) L is oriented upward or downward depending on whether T is red or green. Note that the orientations of R and L

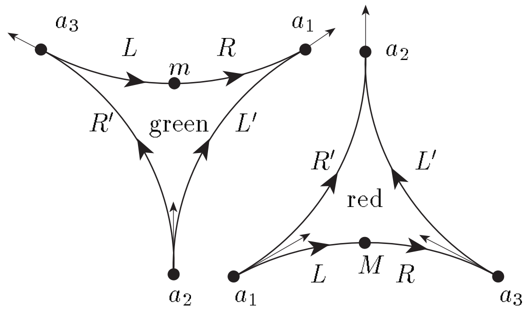

Figure 6: A green (red) pseudotriangle and its canonical orientation.

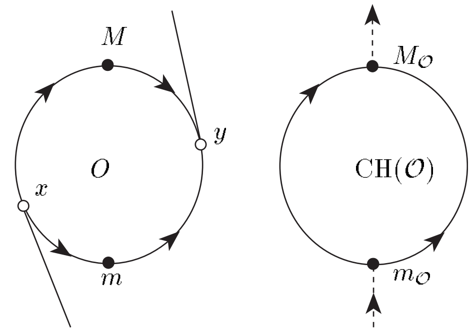

Figure 7: Canonical orientations of an obstacle and (by convention) of the convex hull.

are consistent at the extremal point of T. According to the first part of Lemma 1 one has:

Lemma 2 *The canonical orientation of a pseudotriangle is consistent with the upward orientation of its bitangents.* □

Therefore the canonical orientations of pseudotriangles induce a canonical orientation of $\mathcal{H}_0$. Note that if O is an obstacle with minimal (maximal) point m (M), and if x (y) is the third cusp of the pseudotriangle with extremal point m (M), then the arc xmy (yMx) is oriented counterclockwise (clockwise). It follows that $\mathcal{H}_0$ is acyclic. Similarly the dual directed graph $\mathcal{H}_0^*$ of $\mathcal{H}_0$ (a dual edge is directed from the right side to the left side of the corresponding primal edge) is acyclic. Let $e = (\mathrm{Tail}(e), \mathrm{Head}(e))$ be an edge of $\mathcal{H}_0$ and let $e^* = (\mathrm{Tail}(e^*), \mathrm{Head}(e^*))$ be its dual edge in $\mathcal{H}_0^*$. The accessibility relations in $\mathcal{H}_0$ and $\mathcal{H}_0^*$ are compatible, i.e., the transitive and reflexive closure of the relation defined by $\mathrm{Tail}(e) < e < \mathrm{Head}(e)$ and $\mathrm{Tail}(e^*) < e < \mathrm{Head}(e^*)$ is a partial order on the set of faces, edges, and vertices of $\mathcal{H}_0$.
The complexes X_A. For p a point lying on a bitangent $b \in G_0$ we denote by $\gamma(p)$ the set of upward

rays $\{p\} \times \mathcal{S}^1_+$, and by $\gamma_{\min}(p)$ ($\gamma_{\max}(p)$) the set of upward rays (p,u) that point into the right (left) side of b, respectively denoted by b_R and b_L. We denote by E the *disjoint* union of the closed pseudotriangles in $\mathcal{H}_0$; note that a point-bitangent appears twice in E.

Let A be a down-set of $(G_0, <)$. We denote by J_A ($J_{G_0 \setminus A}$) the endpoints of the bitangents in A ($G_0 \setminus A$). The complex X_A is the 'visibility complex' of the scene whose 'obstacles' are (1) the obstacles in $\mathcal{O}$, (2) the bitangents in A, and (3) the left sides of the bitangents in $G_0 \setminus A$, i.e., b in $G_0 \setminus A$ is an obstacle *only for rays* that pierce b from its left side to its right side. The definition is a last but one variation on quotient space : We form a quotient space V_A of $E \times \mathcal{S}^1_+$ by identifying the rays (p,u) and (q,u) if the pair $((p,u),(q,u))$ belongs to the topological closure of the equivalence relation $\sim_A$ defined by $(p,u) \sim_A (q,u)$ if (1) p and q ly in the interior of E and u is not the direction of a bitangent, (2) p and q are visible along the direction u, (3) $[pq]$ pierces only bitangents in $G_0 \setminus A$, and (4) the bitangents in $G_0 \setminus A$ pierced by $[pq]$ are pierced from right-to-left. The space V_A is locally a two-dimensional set, except

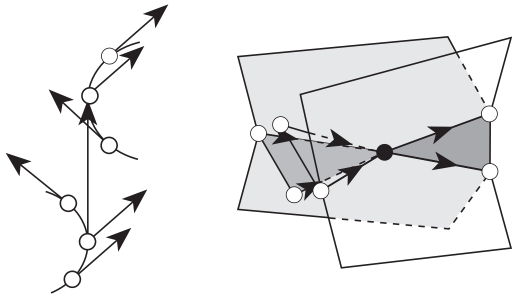

Figure 8: Neighborhood of the visibility complex X_A at a vertex corresponding to a bitangent in $G_0 \setminus A$. Such a vertex is incident to 6 edges and 7 faces.

(1) at (upward) rays $\gamma(\pm\mathcal{O})$, (2) at rays $\gamma(J_A)$, and (3) at rays $\gamma_{\min}(J_{G_0 \setminus A})$. If we fix a direction u in $\mathcal{S}^1_+$ the set of rays in V_A with direction u is locally a one-dimensional set, called the *cross-section* of V_A at u. The curves $\gamma(\pm\mathcal{O})$, $\gamma(J_A)$, $\gamma_{\min}(J_{G_0 \setminus A})$, and the cross-sections at 0 and π induce a 2-dimensional cell decomposition of V_A, denoted by X_A. One can easily check that its vertices are the intersection points of the curves in $\gamma(\pm\mathcal{O})$ (in one-to-one correspondence with the bitangents in $G_0 \cup \Omega(G_0 \setminus A)$) plus the endpoints of the curves $\gamma(J_A)$ and $\gamma_{\min}(J_{G_0 \setminus A})$, that the curves $\gamma_{\min}(J_{G_0})$ and $\gamma_{\max}(J_A)$ are edges, each incident to a unique face. As illustrated in Figure 8, a new type of vertices corresponding to bitangents in $G_0 \setminus A$ appears.

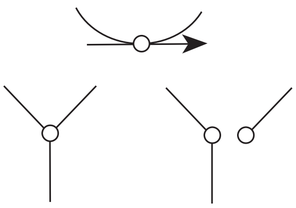

Figure 9: The cutting process.

Planar decomposition of X_A. We introduce now a decomposition of X_A into planar subcomplexes. We cut the visibility complex along each of its edges, lying on the curves γ_i, but we keep glued the two faces incident to the edge that correspond (locally around the edge) to set of rays with the same backward view, as illustrated in Figure 9. In this way we decompose X_A into a set of planar subcomplexes (with pairwise disjoint interiors) X^j_A. The underlying space V^j_A of X^j_A corresponds to the set of upward rays emanating from a convex chain C_j of $\mathcal{H}_0$. These chains are called the canonical chains of X_A. For example consider the case $A = G_0$; the complex X_{G_0} is composed of $2n-2$ connected components : one per pseudotriangle in the pseudo-triangulation. Let $X_{G_0}(T)$ be the visibility complex associated with the pseudotriangle T. If T is red then the complex $X_{G_0}(T)$ is decomposed into 3 planar subcomplexes: namely $r(R', L' \cup L)$ (patches 1 and 2), $r(L \cup R, L' \cup R')$ (patches 3 and 5), and $r(L', R \cup R')$ (patches 4 and 6). The canonical chains are R', $L \cup R$, and L'.

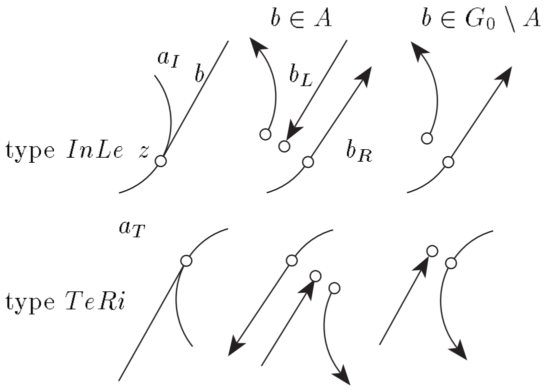

Figure 10: The canonical chain C_j is oriented (this orientation should not be confused with the canonical orientation of $\mathcal{H}_0$.) such that V^j_A is the set of rays pointing on the right side of C_j.

Description of the canonical chains. Let z be an endpoint of a bitangent b in G_0. We denote by b_R, b_L the right and left sides of b, respectively. We denote by a_I, a_T the arc of which z is the initial and terminal point, respectively. The point z is said to be of type In or Te depending on whether z is the initial point or the terminal point of the bitangent b (directed upward). The point z is said to be of type Le or Ri depending on whether the obstacle is on the right side or left side of b. At the beginning of the algorithm both sides of b are obstacles; upon termination only the right side b_R is still an obstacle. The local canonical chains at $b \in G_0$ are described in the following table (see also the above figure).

z type	$b \in$	prefix of $C(b)$	factor	suffix
InLe	A	a_I	$a_T b_R$	b_L
InLe	$G_0 \setminus A$	a_I	$a_T b_R$	
InRi	A	b_R	$b_L a_T$	a_I
InRi	$G_0 \setminus A$	b_R	$a_T a_I$	
TeLe	A	b_L	$b_R a_I$	a_T
TeLe	$G_0 \setminus A$		$b_R a_I$	a_T
TeRi	A	a_T	$a_I b_L$	b_R
TeRi	$G_0 \setminus A$		$a_I a_T$	b_R

Figure 11: Local canonical chains at $b \in G_0$.

Update of (the canonical decomposition of) X_A. Let $b \in \max_< A$. We explain now how to compute $X_{A\setminus\{b\}}$ from X_A.

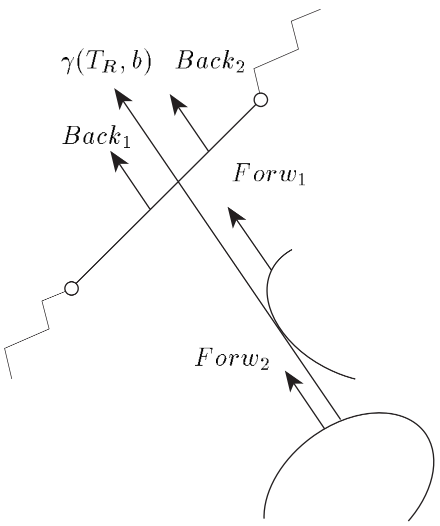

Figure 12: Update of X_A.

Let T_R (T_L) be the pseudotriangle incident upon b along its right (left) side. We denote by $\gamma(T_R, b)$ the curve of rays emanating from T_R, tangent to T_R, and with forward view the (right) side of b. This curve is represented by the sequence $\Omega(b)$.

Let $Backward(b_L)$ $(Forward(b_R))$ be the set of rays in V_A with backward (forward) view the left (right) side b_L (b_R) of b. Clearly the space $V_{A\setminus\{b\}}$ is obtained from V_A by identifying rays in $Backward(b_L)$ and in $Forward(b_R)$ that are supported by the same line.

Let $C = C_1 b_L C_2$ be the canonical chain of X_A that contains the left side b_L of b. The set $Backward(b_L)$ is a subset, bounded by the curves $\gamma_{\max}(\partial b)$, of (the underlying set of) the planar complex $X_A(C)$. The curve of rays $\gamma(T_R, b)$ splits $Backward(b_L)$ into two parts denoted by $Back_1(b_L)$ and $Back_2(b_L)$.

The curve of rays $\gamma(T, b)$ splits $Forward(b_R)$ into two parts $Forw_1(b_R)$ and $Forw_2(b_R)$. Part $Forw_i(b_R)$ is a subset, bounded by the curves $\gamma_{\max}(\partial b)$, of the (underlying space of) a canonical subcomplex of X_A, denoted by $X_A(C_i')$.

Upon removal of b the left side b_L of b disappears from the canonical chains and we should rearrange the chains C_1, C_2, C_1' and C_2' to create new canonical chains. The patches $Back_i(b_L)$ and $Forw_i(b_R)$ are then used to update their corresponding canonical complexes. We distinguish several cases.

Case 1. The endpoints of b are cusp points of T_L. In that case $C = b_L$ and $Backward(b_L) = |X_A(C)|$. The chain C disappears and we should just update the the complexes associated with the canonical chains C_i'. This is done as follows. **1**–We introduce in $X_A(C_i')$ the curves $\gamma_{\max}(\partial b)$. These curves bound the single patch $Forw_i(b_R)$ in $X_A(C_i')$. **2**– We cut the complex $X_A(C)$ along $\gamma(T_R, b)$. (This is done in time proportional to the number of bitangents in $\Omega(b)$ with the representation of B introduced in subsection 2.) The resulting piece $Back_i(b_L)$ is glued with the piece $Forw_i(b_R)$ along their common boundaries supported by the curves $\gamma_{\max}(\partial b)$ and $\gamma(T_R, b)$. **3**– We remove the patches $Forw_i(b_R)$ and the curves $\gamma_{\max}(\partial b)$.

Case 2. The endpoints of b are cusp points of T_R. In that case the curves $\gamma_{\max}(\partial b)$ appears in the boundary of $X(C_i')$. The chains C_i and C_i' are concatenated to create a new chain C_i''. To create $X_A(C_i'')$ we proceed as follows **1**– We cut the complex $X_A(C)$ along $\gamma(T_R, b) = \gamma(T_R)$. The resulting piece that contains $Back_i(b_L)$ is glued with $X_A(C_i')$ along $\gamma(T_R, b)$; **2**– The piece $Forw_i(b_R)$ is removed.

Case 3 and 4. One endpoint of b is a cusp point of T_R and the other is a cusp point of T_L. Similar to the two previous cases.

From the above discussion we get:

Lemma 3 *The set of planar subcomplexes associated with $X_{A\setminus\{b\}}$ can be computed from the set of planar subcomplexes associated with X_A is time proportional to the number of bitangents in $\Omega(b)$.*

Keeping the history of the construction of the sequence of complex X_A we get the ray-shooting data structure in time $O(k)$ up to some polylog factor, due to the use of dynamic point location data structures.

4 Covering problems.

According to L. Fejes Tóth [5], the boundary of a set of $n \geq 3$ interior pairwise disjoint convex sets can be illuminated by $4n-7$ points. The proof of L. Fejes Tóth proceeds by growing the convex sets unboundedly in all directions but the growth (in a given direction) is limited by the condition that the convex sets remain pairwise interior disjoint. In this way the convex sets will expand into convex polygons that fill the plane except for a finite number of gaps that are also convex polygons. A suitable choice of the lighting points at the vertices of the gaps leads to the $4n-7$ bound. It is not clear how to turn this 'growing process' into an (efficient) algorithm to compute a lighting set. L. Fejes Tóth provides also sets of $n \geq 3$ convex sets which cannot be illuminated by less than $4n-7$ points but leaves open a practical characterization of all cases when this number of lighting points is required. It turns out that the computation and characterization problems can be solved using the concept of pseudo-triangulation. Let us say that a visibility complex requires x lighting points if x lighting points are always sufficient and sometimes necessary to illuminate the boundary of any realization of the visibility complex.

Theorem 3 *Computing a lighting set for a set of n pairwise disjoint convex sets reduces in $O(n)$ time to computing a pseudo-triangulation.*

The visibility complexes requiring $4n-7$ lighting points are in one-to-one correspondence with the triangular planar graphs on n vertices.

Proof. Since the boundary points of the convex sets are the boundary points of the pseudotriangles it is sufficient to find a lighting set for the pseudotriangles of a pseudo-triangulation. How many lighting points are necessary for a pseudotriangle? Two in general (computable in $O(1)$ time) but only one if a side of the pseudotriangle reduces to a line segment. Let a be the number of pseudotriangles with a line segment side or, equivalently, the number of exterior[6] bitangents in the pseudo-triangulation. From the above discussion a lighting set with $4n-4-a$ lighting points exists and is computable in $O(n)$ time from the pseudo-triangulation. It is no hard to see that $a \geq 3$, from which the result follows (taking care to send the lighting points on the convex hull far enough to illuminate also the boundary of the convex hull).

According to the previous discussion a necessary condition for a visibility complex to require $4n-7$ points is that no more than three free exterior bitangents exist. This condition implies strong conditions on the visibility complex: (1) the number of exterior bitangents is three; they all ly on the convex hull; (2) the visibility complex depends only on the planar graph whose vertices are the obstacles and whose edges are the interior bitangents of any pseudo-triangulation. From (1) we can deduce that this planar graph is triangular. Conversely, according to the Koebe Representation Theorem' [9, page 96], any triangular planar graph on n points is realizable as the contact graph of a set of n interior disjoint circles. While configurations of circles require only $2n-2$ lighting points (see [5]), a slight perturbation of the circles leads to configurations that require $4n-7$ lighting points. □

[6] A bitangent to obstacles O and O' is said to be interior (exterior) if its supporting line separates (does not separate) the obstacles O and O'.

A polygonal *cover* of a set $\{O_i\}$ of n pairwise disjoint convex sets is a set $\{O'_i\}$ of pairwise disjoints convex polygons such that $O_i \subseteq O'_i$. H. Edelsbrunner et *al* [6] have shown that no more than $6n-9$ sides and $3n-6$ slopes for $n \geq 3$ are required for a polygonal cover—these bounds being optimal in the worst case; the proof is based on a growing process similar to the one used by L. Fejes Tóth, and therefore doesn't lead to an (efficient) algorithm. Once more the notion of pseudo-triangulation is the key idea to achieve optimal time complexity.

Theorem 4 *Computing a worst case optimal (with respect to the number of slopes and sides) polygonal cover of a set of n pairwise disjoint convex sets reduces in $O(n)$ time to computing a pseudo-triangulation.*

Proof. Let $\mathcal{D}$ be the set of closed half-planes bounded by the supporting lines of the bitangents of a pseudo-triangulation. Let $\mathcal{D}_i$ be the set of the half-planes D in $\mathcal{D}$ such that (1) $O_i \subseteq D$ and (2) the bitangent that defines D is tangent to O_i. Since O_i is convex, $\mathcal{D}_i$ is a non-redundant presentation of the convex polygon $P_i = \bigcap \mathcal{D}_i$. Clearly $\{P_i\}$ is a convex polygonal cover with $6n-6$ sides realizing no more than $3n-3$ slopes. The computation of $\{P_i\}$ reduces clearly in linear time to the computation of the pseudo-triangulation. □

5 Realizability questions

We say that an arrangement of n pseudolines is k-*stretchable* if it is isomorphic to an arrangement of pseudolines which satisfies the following property: its number of transversal intersection points with any (straight) line doesn't exceed $2k+n$. Clearly an arrangement of pseudolines is stretchable iff. it is 0-stretchable. Similarly an arrangement of pseudolines is 1-stretchable iff. it is realizable by pairwise disjoint pseudotriangles. Arrangements of pseudolines realizable by $k+1$ by $k+1$ disjoint pseudotriangles are examples of k-stretchable arrangements.

Theorem 5 *Any arrangement of n pseudolines is realizable by set of n pseudotriangles. (Therefore any arrangement of n pseudolines is n-stretchable).*

Proof. Omitted from this version. □

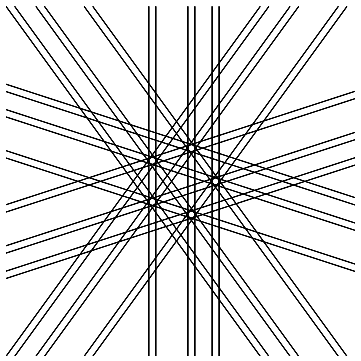

Figure 13: The regular k–gon, with small circles at its vertices.

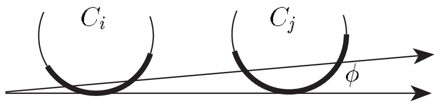

Figure 14: Perturbing $L(C_i, C_j) \in L_h$, for $\{i,j\} \in E$. Here $\tau_{(E,\sigma)}(i,j,-h) = \tau_{(E,\sigma)}(i,j,h) = \sigma(\{i,j\}) = +1$.

Theorem 6 *The number of arrangements of n pseudolines realizable by disjoint pseudotriangles (= 1-stretchable) is $2^{\Theta(n^2)}$.*

To prove this theorem we are going to define a class of $2^{n^2/8}$ sets[7] of n pairwise disjoint convex obstacles such that the arrangements of the dual curves of this convex obstacles (called the *order types* of the sets, by analogy with the order types of sets of points) are all combinatorially distinct, and therefore in number $2^{n^2/8}$. Consider a pseudo-triangulation of any of these sets. The dual image of its pseudotriangles is an arrangement of pseudolines, that is obviously realizable by *disjoint* pseudotriangles. Since there are only $2^{O(n \log n)}$ different pseudo-triangulations for each set, our result follows.

[7]Related to an example in [10].

Theorem 7 *The number of order types of sets of n disjoint convex objects in the plane is at least*
1. $2^{n^2/8}$*;*
2. $2^{\Omega(dn \log n)}$*, if the objects are of degree $d = O(n^\alpha)$, for some fixed α with $0 \le \alpha < 1$.*

In the proof we need the following lemma, whose (not very difficult) proof we omit from this version.

Lemma 4 *The number of labeled graphs with n vertices and maximal degree at most d is $2^{\Omega(dn \log n)}$, provided $d = O(n^\alpha)$, for some fixed α with $0 \le \alpha < 1$.*

Remark For an asymptotically sharp result, under stricter conditions on d, we refer to [11]. Since the proof of the latter result is quite involved, we prefer to give the simple proof of the weaker lemma 4.

Proof. We shall prove both parts simultaneously. Let $k = \lceil n/2 \rceil$. Consider a regular k–gon with vertices $p_1, \ldots, p_k$. Put a small circle C_i of radius ϱ centered at p_i. Here ϱ is small enough to guarantee that no line in the plane intersects more than 2 of the circles. We denote by $L(\epsilon C, \epsilon' C')$, with $\epsilon, \epsilon' \in \{+,-\}$, the tangent line that is directed from C to C' and contains the objects C and C' in their left or right half-planes according to the sign ϵ and ϵ' in front of C and C'. Draw all common tangent lines of any pair of circles parallel to the sides and the diagonals of the k-gon (so for any pair of distinct circles we draw exactly 2 out of their 4 common tangent lines), see Figure 13. This set of lines is partitioned into k classes of parallel lines, denoted by $L_{k+1}, \ldots, L_{2k}$. All lines in class L_h are given the same, arbitrarily chosen, direction. Let $\mathcal{T}$ be the set of triples (i,j,h) such that $L(C_i, C_j) \in L_{|h|}$. So for $(i,j,h) \in \mathcal{T}$ we have $1 \le i, j \le k <| h |\le 2k$, $i \ne j$, and $(i,j,h) \in \mathcal{T}$ iff. $(i,j,-h) \in \mathcal{T}$. Obviously $|\mathcal{T}| = \Theta(k^2)$. Let $V = \{1, \ldots, k\}$, and let $\mathcal{G}$ be the set of labeled undirected graphs, in the first case, and the set of labeled undirected graphs of maximal degree not exceeding $d-2$ in the second case. The restriction on the degree will become clear from the construction below. For a pair (E, σ), such that $(V, E) \in \mathcal{G}$ and $\sigma : E \to \{-1, +1\}$, define $\tau_{(E,\sigma)} : \mathcal{T} \to \{-1, +1\}$ by

$$\tau_{(E,\sigma)}(i,j,h) = \begin{cases} \sigma(\{i,j\}), \text{if } \{i,j\} \in E, \\ -1, \text{if } \{i,j\} \notin E \text{ and } h > 0, \\ +1, \text{if } \{i,j\} \notin E \text{ and } h < 0. \end{cases}$$

Note that $\tau_{(E,\sigma)} \ne \tau_{(E',\sigma')}$ for $(E,\sigma) \ne (E',\sigma')$. Hence there are at least $\sum_{E:(V,E)\in\mathcal{G}} 2^{|E|} \geq |\mathcal{G}|$ such mappings $\mathcal{T} \to \{-1,+1\}$. Since the number of graphs in $\mathcal{G}$ is $2^{\binom{k}{2}}$, in the first case, and $2^{\Omega(dk\log k)}$ in the second case, the proof is complete, provided we show that every $\tau_{(E,\sigma)}$ is *realizable*. By this we mean that there is a set of $2k$ disjoint convex objects $O_1,\ldots,O_{2k}$ such that $\tau_{(E,\sigma)(i,j,\pm h)}$ satisfies, for all triples $(i,j,\pm h) \in \mathcal{T}$ with $h > 0$:

condition ($\star$): $\tau_{(E,\sigma)}(i,j,\pm h) = 1(-1)$ if the support line of $\pm O_h$, parallel to $L(O_i,O_j)$, lies to the left (right) of $L(O_i,O_j)$.

(By convention a support line of O_h $(-O_h)$ contains O_h in its left (right) half plane.)

So let us describe the construction of the convex objects $O_1,\ldots,O_{2k}$ for some fixed $\sigma : E \to \{-1,+1\}$. These objects are obtained by

1. slightly perturbing the objects bounded by the circles C_i; this yields objects O_i, $1 \leq i \leq k$;
2. adding a convex object O_h, $k < h \leq 2k$, that intersects all lines in L_h ahead of the regular k–gon.

On each circle C_i, $1 \leq i \leq k$, we introduce a set T_i of $2k$ disjoint small chords, centered at the points whose tangent lines are parallel to the sides and diagonals of the k-gon. (See Figure 14.)

Consider a triple $(i,j,h) \in \mathcal{T}$ with $h > 0$. Note that $L(C_i,C_j) \in L_h$. If $\{i,j\} \in E$ we perturb the line $L(C_i,C_j)$ into a line $L_{(E,\sigma)}(i,j)$, such that

(i) the tilt of $L_{(E,\sigma)}(i,j)$ with respect to $L(C_i,C_j)$ is $\pm\varphi$ if $\sigma(i,j) = \mp 1$;

(ii) $L_{(E,\sigma)}(i,j)$ intersects C_i (C_j) in the same chord of T_i (T_j) as $L(C_i,C_j)$.

It is not hard to see that there is a small $\varphi > 0$ satisfying these conditions. (See also Figure 14.) The sign of the tilt is determined by our intention to insert a convex object O_h that intersects $L(C_i,C_j)$ ahead of C_i and C_j, and that lies to the right (left) of $L_{(E,\sigma)}(i,j)$ iff. $\tau_{(E,\sigma)}(i,j,h) = -1$ $(+1)$. If $\{i,j\} \notin E$ we take $L_{(E,\sigma)}(i,j) = L(C_i,C_j)$. Note that this way we perturb exactly d_i of the lines tangent at C_i, where d_i is the degree of $i \in V$ in the graph (V,E).

We first put, for $1 \leq i \leq k$, a convex object O_i at vertex p_i of the regular k–gon, that is tangent to all lines of the form $L_{(E,\sigma)}(i,j)$ or $L_{(E,\sigma)}(j,i)$. To this end consider the convex object O'_i bounded by the circle C_i, and the d_i lines $L_{(E,\sigma)}(i,j)$ or $L_{(E,\sigma)}(j,i)$, that intersect the interior of this circle. Note that we still have to perturb object O'_i so that its boundary becomes algebraic (and of degree at most d in case 2). However, all of the k lines of the form $L_{(E,\sigma)}(i,j)$ or $L_{(E,\sigma)}(j,i)$ are tangent to it.

We now introduce a convex object O_h (of degree $O(1)$) that

(i) intersects all lines $L(C_i,C_j) \in L_h$ ahead of C_i and C_j;

(ii) lies to the right (left) of $L_{(E,\sigma)}(i,j)$ iff. $\tau_{(E,\sigma)}(i,j,h) = -1$ $(+1)$.

Condition (i) implies that $\tau_{(E,\sigma)}(i,j,\pm h)$ satisfies condition ($\star$), if $L_{(E,\sigma)}(i,j) \neq L(C_i,C_j)$, viz. if $\{i,j\} \notin E$. Condition (ii) implies that $\tau_{(E,\sigma)}(i,j,\pm h)$ satisfies condition ($\star$), if $L_{(E,\sigma)}(i,j) = L(C_i,C_j)$, viz. if $\{i,j\} \in E$. Therefore the order type of the set $\{O'_1,\ldots,O'_k,O_{k+1},\ldots,O_{2k}\}$ is a realization of $\tau_{(E,\sigma)}$. □

In the full version of the paper we show that all sets in the proof of theorem 7 have distinct visibility graphs/complexes. This shows:

Corollary 1 *The number of visibility graphs/complexes of sets of n disjoint convex objects in the plane is at least*
1. $2^{n^2/8}$;
2. $2^{\Omega(dn\log n)}$ *if the objects are of degree* $d = O(n^\alpha)$, *for some fixed* α *with* $0 \leq \alpha < 1$.

References

[1] M. Pocchiola and G. Vegter. The visibility complex. *Internat. J. Comput. Geom. Appl.*, 1996. To appear in the special issue of Internat. J. Comput. Geom. Appl. devoted to the 9th Annu. ACM Sympos. Comput. Geom. held in San Diego in June 1993.

[2] M. Pocchiola and G. Vegter. Computing visibility graphs via pseudo-triangulations. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 248–257, June 1995.

[3] M. de Berg. *Ray Shooting, Depth Orders and Hidden Surface Removal*, volume 703 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1993.

[4] B. Chazelle. *Computational geometry: a retrospective*, volume 1 of *Lecture Notes Series on Computing*, pages 22–46. World Scientific, 2 edition, 1995.

[5] L. Fejes Tóth. Illumination of convex discs. *Acta Math. Acad. Sci. Hungar.*, 29(3–4):355–360, 1977.

[6] H. Edelsbrunner, A. D. Robison, and X. Shen. Covering convex sets with non-overlapping polygons. *Discrete Math.*, 81:153–164, 1990.

[7] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, and G. Ziegler. *Oriented Matroids*. Encyclopedia of Mathematics and Applications. Cambridge University Press, 1993.

[8] J. Richter-Gebert and G. M. Ziegler. Realization spaces of 4-polytopes are universal. *Bulletin of the American Mathematical Society*, 32(4), October 1995.

[9] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, New York, NY, 1995.

[10] J.E. Goodman and R. Pollack. Multidimensional Sorting. *Siam J. Comput.*, 12:484–507, 1983.

[11] E. A. Bender, E. R. Candfield, and B. D. McKay. The asymptotic number of labeled connected graphs with a given number of vertices and edges. *Random Structures and Algorithms*, 1(2):127–169, 1990.