The Visibility Complex

(Extended Abstract)

Michel Pocchiola* Gert Vegter†

Abstract

We introduce the visibility complex of a collection $\mathcal{O}$ of n pairwise disjoint convex objects in the plane. This 2–dimensional cell complex may be considered as a generalization of the tangent visibility graph of $\mathcal{O}$. Its space complexity k is proportional to the size of the tangent visibility graph. We give an $O(n \log n + k)$ algorithm for its construction. Furthermore we show how the visibility complex can be used to compute the view from a point or a convex object with respect to $\mathcal{O}$ in $O(m \log n)$ time, where m is the size of the view. The view from a point is a generalization of the visibility polygon of that point with respect to $\mathcal{O}$.

1 Introduction

Consider a collection $\mathcal{O}$ of pairwise disjoint objects in the plane. We are interested in problems in which these objects arise as obstacles, either in connection with visibility problems where they can block the view from an other geometric object, or in motion planning, where these objects may prevent a moving object from moving along a straight line path. The visibility graph is a central object in the context of these problems. For polygonal obstacles the vertices of these polygons are the nodes of the visibility graph, and two nodes are connected by an arc if the corresponding vertices can see each other. [Wel85] describes the first non-trivial algorithm for computing the visibility graph of a polygonal scene with a total of n vertices in $O(n^2)$ time. [GM91] presents an optimal $O(n \log n + k)$ algorithm, where k is the number of arcs of the visibility graph. A related problem concerns the computation of the view (visibility polygon) of a point amidst polygonal obstacles. There are several output sensitive algorithms for a single shot computation, see [?] and the references given there.

Due to its discrete structure the visibility graph is not rich enough to maintain the view of a moving point in a continuously varying direction. To cope with this and similar problems we introduce the *visibility complex* of a set of pairwise disjoint convex objects $\mathcal{O}$, a 2–dimensional cell complex that can be considered as a subdivision of the set of rays emanating from these objects. Faces correspond to collections of rays of 'constant visibility'.

Similar ideas have been used in earlier work on visibility, shortest paths and motion planning amidst polygonal obstacles, see e.g. [CG89, Poc90, Veg90, Veg91]. Here the space of directed lines, endowed with a partition generated by the set of obstacles, is regarded as the main structure, instead of the scene of obstacles itself.

Unless otherwise stated we assume that the convex objects have complexity $O(1)$, so we can compute the tangents from a point to a convex object, as well as the common tangents of two objects in $O(1)$ time. Then the space complexity of the visibility complex is proportional to the size of the *tangent visibility graph (TVG)* of $\mathcal{O}$. The set of vertices of the latter graph is $\mathcal{O}$. Furthermore any common tangent of two objects $O_1, O_2 \in \mathcal{O}$ whose endpoints can see each other correspond to an edge $\{O_1, O_2\}$ of the TVG. (Note that there are at most 4 edges between two vertices.) We show that the visibility complex can be computed in optimal $O(n \log n + k)$ time. Here k is the complexity of the visibility complex, or, equivalently, the number of arcs of the tangent visibility graph.

*e-mail: pocchiol@dmi.ens.fr
†e-mail: gert@cs.rug.nl

0

The visibility complex contains sufficient information to maintain the view along a moving ray. Such a continuously moving ray corresponds to a curve in the visibility complex. Positions of the ray at which the visibility changes correspond to intersections of the curve with edges of the visibility complex. In particular if the moving ray rotates around its origin, maintaining the visibility boils down to computing the *visibility polygon* (or: the view) of this origin. We show how to compute the view from a point in $O(h \log n)$ time, where h is the size of the view. This is the main result of section 2.

2 The Visibility Complex

Terminology

First we introduce some terminology. Consider a collection $\mathcal{O}$ of pairwise disjoint convex obstacles. Unless otherwise stated each obstacle is *strictly* convex, and has a *smooth* boundary. As mentioned in the introduction the complexity of each object is $O(1)$. For the sake of convenience we assume that the objects in $\mathcal{O}$ are in *general position*, in the sense that no three objects share a common tangent line. To facilitate a uniform description of the visibility complex we introduce an object O_∞ at infinity, which can be viewed as a sufficiently large circle that encloses the collection of obstacles. The complement $\mathcal{F}$ of the union of the set of obstacles in the disc enclosed by O_∞ will be called *free space*.

Any finite sequence of points on the boundary of a convex object O subdivides the boundary into curved segments, called *arcs*. A *bitangent* is a free line segment that is tangent to two objects at its endpoints. A *chain* is a simple curve consisting of an alternating sequence $s_1, \cdots, s_m$ of bitangents and arcs, such that s_i and s_{i+1} share an endpoint, at which the bitangent is tangent to the arc. Such a chain is called *convex* if connecting its endpoints by a line segment yields a simple closed curve that bounds a convex region. A *maximal (minimal) point* of a convex chain is a boundary point at which the tangent line to the boundary is horizontal, such that the chain lies *below (above)* this tangent line. An *extremal point* is either a maximal or a minimal point. A *pseudo-triangle* is a simply connected subset R of $\mathcal{F}$ such that (i) the boundary ∂R is a sequence of three convex chains, that are tangent at their endpoints, and (ii) R is contained in the triangle formed by the three endpoints of these convex chains (also see Figure 7).

The underlying space

For a point $p \in \mathcal{F}$ we are interested in the object that we can see from p in a certain direction $u \in \mathbb{S}^1$. Note that this *view from* p in the direction u is the same as the view from p_u in the direction u, where $p_u \in \partial\mathcal{F}$ is the first obstacle point that is hit when moving from p in the direction $-u$. So if we are able to determine p_u for any pair $(p, u) \in \mathcal{F} \times \mathbb{S}^1$ it suffices to know the view from p_u in the direction u. Furthermore the point we see from p is of the form $p + ru$, for some positive scalar r. If we see a point on the object at infinity we have $r = \infty$.

These simple observations motivate the following more formal definition. Let $V_0 \subset \partial\mathcal{F} \times \mathbb{S}^1 \times \mathbb{R}$ be the set defined by $(x, u, r) \in V_0$ iff. (i) $r > 0$ and (ii) $x, x + ru \in \partial\mathcal{F}$ and (iii) $(x, x + ru) \subset \mathcal{F}$. The closure of V_0 as a subspace of $\mathbb{R}^2 \times \mathbb{S}^1 \times \mathbb{R}$ is denoted by V_1. With $\xi = (x, u, r) \in V_1$ we associate the closed line segment $\text{seg}(\xi)$ defined by $\text{seg}(\xi) = [x, x + ru]$. Note that for a pair $(x, u) \in \partial\mathcal{F} \times \mathbb{S}^1$ there is at most one positive $r \in \mathbb{R}$ such that $(x, u, r) \in V_0$. Therefore V_0, and hence V_1, is a 2–dimensional set.

Recall that one of our goals is to maintain the view from a continuously moving point $p(t)$ in a continuously changing direction $u(t)$, i.e. we are interested in the view associated with a continuous curve $\gamma : t \mapsto (p(t), u(t))$. Note that with γ we can associate—more or less naturally—the curve $\overline{\gamma} : t \mapsto (p_{u(t)}(t), u(t), r(t)) \in V_1$, such that the view from $p(t)$ in the direction $u(t)$ is the point $p_{u(t)}(t) + r(t)u(t)$. The curve $\overline{\gamma}$ is *not continuous* in general. Consider e.g. figure 1, where $p_u(t)$ ranges over the curve γ_0 in the plane, and $u(t) = u_0 \in \mathbb{S}^1$ is the vertically upward direction. Here $t \mapsto r(t)$ is discontinuous at positions where the line $p(t) + \mathbb{R}u_0$ is tangent to an obstacle. Furthermore $p_{u(t)}$ is discontinuous when $p(t) + \mathbb{R}u_0$ is tangent to O_1.

To associate a continuous curve with γ we identify certain points of V_1. More precizely for $\xi_1, \xi_2 \in V_1$, with $\xi_i = (x_i, u_i, r_i)$, we say that $\xi_1 \equiv \xi_2$ iff. $r_1, r_2 > 0$ and $u_1 = u_2$ and $\text{seg}(\xi_1) \subset \text{seg}(\xi_2)$ or $\text{seg}(\xi_2) \subset \text{seg}(\xi_1)$. The transitive closure of this relation is an equivalence relation on V_1, which we again denote by $\equiv$. Finally V is the quotient space of V_1 with respect to $\equiv$, endowed with the quotient topology. The set V is the *underlying set of the visibility complex*. If we fix a direction $u_0 \in \mathbb{S}^1$ the set $V \cap \{u = u_0\}$ is locally a one-dimensional set (except at a finite number of points). We shall refer to this set as the *cross section* of the visibility complex at $u = u_0$. A representation of this set for a configuration of three obstacles is depicted in Figure 1 below. Note that the image of the curve $\overline{\gamma}$ under the quotient map $q : V_1 \to V$ is a *continuous* curve in V.

The combinatorial structure

We shall now turn V into a two-dimensional cell–

Figure 1: a. A configuration of three obstacles in the plane; b. The set $V \cap \{u = u_0\}$, where $u_0 \in \mathbb{S}^1$ is directed vertically upward.

complex. The corresponding incidence structure will be the basis for our choice of a data structure representing the visibility complex.

Let $q : V_1 \to V$ be the quotient map as defined in the previous subsection. A *face* (*edge, vertex*) is a connected component of the set of points $x \in V$ for which the number of points in $q^{-1}(x)$ is equal to 1 (3, more than 3, respectively). Note that an edge corresponds to a set of line segments whose endpoints are on obstacle boundaries, whereas the segments are tangent to the *same* obstacle. Similarly a vertex corresponds to a line segment that is tangent to two obstacles.

If the obstacles are in general position (as we assume in this paper) every edge is incident to three faces, and two vertices. Furthermore every vertex is incident to four edges and six faces. To see this consider Figure 2, where we depict the topology of the visibility complex near a vertex corresponding to a bitangent of two objects O_1 and O_2. Let θ_0 be its slope, and let $u_0 \in \mathbb{S}^1$ ($u_0^{\pm} \in \mathbb{S}^1$) have slope θ_0 ($\theta_0 \pm \epsilon$, for some small positive ϵ). It is not hard to assemble the cross–sections of the visibility complex at $u = u_0^-$, $u = u_0$ and u_0^+ into the configuration of 4 edges and 6 faces near the vertex.

Definition 1 *The two–dimensional cell–complex defined above is called the visibility complex of the set of obstacles $\mathcal{O}$.*

The planar subcomplex of an obstacle

Figure 2: A neighborhood of a vertex of the visibility complex.

With each obstacle O in $\mathcal{O} \cup \{O_\infty\}$ we associate a subcomplex in the following way.

Consider the set $V_0(O) = \{(x, u, r) \in V_0 \mid x \in \partial O\}$, which is a closed subset of V_0. In fact this subset may be identified with the set of rays emanating from the boundary of O and pointing into free space. Therefore the quotient map $q : V_1 \mapsto V$ maps it onto a subset $V(O)$ of V, that has the structure of a *planar subcomplex* of the visibility complex, which we shall denote by $\mathcal{P}(O)$.

In the situation of figure 2 the faces labeled 1, 2 and 3 belong to the subcomplex $\mathcal{P}(O_0)$, those labeled 5 and 6 to $\mathcal{P}(O_1)$ and face 4 belongs to $\mathcal{P}(O_2)$. Note that each edge 'belongs' to two subcomplexes, and each vertex to three.

To describe the subcomplex $\mathcal{P}(O)$in more detail (and to see that is planar) we shall endow $V(O)$ with global coordinates, thereby mapping it onto a subset of the plane.

First recall that a convenient parametrization of the set of directed lines in the plane is given by the *polar coordinates of a directed line*: we identify this set of lines with the cylinder $[0, 2\pi) \times \mathbb{R}$ using the bijection which maps the pair (θ, u) on the directed line $y \cos\theta - x \sin\theta - u = 0$ with slope θ and signed distance u to the origin.

Since there is a 1:1–correspondence between the set of rays emanating from ∂O that point into free space and the set $\mathcal{L}_O$ of directed lines intersecting O, we obtain global coordinates on $V(O)$ by passing to the polar coordinates on $\mathcal{L}_O$.

Convention In the sequel we restrict to the set of lines whose slope lies in the range $[0, \pi]$, unless stated otherwise. So non–horizontal lines will be directed upward. For a point a in the plane the horizontal lines through a with slopes 0 and π will be denoted by h_a and $\overline{h}_a$, respectively.

Example 2 Consider a convex object O with minimal point m_0 and maximal point m_1. The region corresponding to the set of lines emanating from this object and pointing into free space is depicted in Figure 3a. The upper boundary of this region is a curve

Figure 3: The set of directed lines intersecting (a) a convex object and (b) bounded segment.

with endpoints h_{m_1} and $\overline{h}_{m_0}$.

Example 3 Consider a face in the subcomplex $\mathcal{P}(O_\infty)$, see Figure 4. It is bounded by edges e_i, $1 \leq i \leq 5$, each corresponding to a set of directed lines tangent to O_i. The *minimal (maximal) vertex* l_{min} (l_{max}) corresponds to a common tangent of O_5 and O_4 (O_3 and O_4). In the example two successive edges correspond to lines tangent to distinct objects. However this need not be true if the shaded region contains more than one object visible from O_4. If there are m tangent visibilities in this region, then edge e_4, incident upon l_{max}, is subdivided into $m+1$ subedges. We shall call the union of these edges (e_4 in our example) a *fat* edge. During traversal of the visibility complex special care must be taken if we pass a fat edge, see section 3.

The minimal and maximal vertices subdivide the boundary of a face into two chains, called the *right* and *left* boundary of this face. In our example these chains are e_1, e_2, e_3 and e_4, e_5, respectively.

Figure 4: A face of $\mathcal{P}(O_\infty)$.

In the sequel we shall represent the visibility complex by the collection of planar subcomplexes $\mathcal{P}(O)$, $O \in \mathcal{O} \cup \{O_\infty\}$, where each edge is augmented with a pointer to the other edge with which it has to be identified. In this way we can access the faces and vertices incident upon a given edge in $O(1)$ time.

Extension: convex chains

We will often deal with sets of rays emanating from a (convex) chain. To determine the subregion corresponding to the set of rays emanating from the convex chain with endpoints a and b, and pointing in to free space, note that the set of upward directed rays through a emanating from the object is a curve $\mathcal{L}_a$ in this region. It connects the tangent line t_a at a with h_a or $\overline{h}_a$, depending on whether the object lies to the left (as in Figure 3a) or to the right of t_a. The curves $\mathcal{L}_a$ and $\mathcal{L}_b$ bound a region (shaded in Figure 3a) corresponding to the set of lines intersecting the chain am_1b.

As a special case consider the set of lines intersecting a straight line segment pq supported by an upwardly directed line t. This set consists of two regions, corresponding to the set of lines intersecting pq from left to right and from right to left, respectively, see Figure 3b.

Example 4 Consider a pseudo–triangle $a_0a_1a_2$, see Figure 5. There is a well–defined tangent ray t_i at a_i. Let σ_i be the side opposite vertex a_i. The set of rays emanating from side σ_2, pointing into the interior of the pseudo-triangle, is depicted in the rightmost part

Figure 5: A pseudo–triangle and the set of rays emanating from a side (σ_2) pointing into its interior.

of Figure 5. Note that it is similar to the shaded region in Figure 3a. It is however subdivided into two parts by the set of tangents to the sides σ_0 and σ_1. These parts correspond to the sets of rays, emanating from σ_2, along which we can see σ_0 or σ_1, as indicated by the labels in Figure 5.

Remark 5 The curve $\mathcal{L}_a$, introduced above, will be called the *canonical image* of a. The canonical image of a curve is the set of its directed tangent lines.

3 Computing views

We show how the visibility complex can be used to compute the view from a free convex object γ with respect to the set of obstacles. If the obstacles are polygons and the object is a point this amounts to computing the visibility polygon of the point. For simplicity we only consider the case in which the object γ lies outside the convex hull of the set of obstacles. We refer to the full paper for general results. We shall compute the tangencies in the view of γ in polar order, starting with the counterclockwise tangent t of γ with slope 0. When l ranges over the set of all counterclockwise tangents of γ it describes a curve $\overline{\gamma}$ in the subcomplex $\mathcal{P}(O_\infty)$. If γ lies outside the convex hull of $\mathcal{O}$ then $\overline{\gamma}$ intersects the boundary of a face in $\mathcal{P}(O_\infty)$ in at most two points. The sequence of intersections of $\overline{\gamma}$ with edges of $\mathcal{P}(O_\infty)$ corresponds to the view of γ.

Using a trivial auxiliary data structure we can compute the first tangency of the view in $O(\log n)$ time. So suppose t is the current intersection of $\overline{\gamma}$ with an edge of $\mathcal{P}(O_\infty)$, lying in the right boundary of face f. A simple way to find the other intersection of $\overline{\gamma}$ with the (left) boundary of f is to do a binary search on the sequence of edges in the boundary of f. This takes $O(\log n)$ time for each intersection. We refer to this approach as *crossing faces*.

We can also find the next intersection by walking along the left boundary starting at the minimal vertex, until we find an edge containing a line l' that is tangent to γ. Obviously l' is the next intersection of $\overline{\gamma}$ with an edge of $\mathcal{P}(O_\infty)$. The sequence of edges traversed is called the *visible zone* of γ. Although we may need to spend $O(n)$ time to find a single tangency in the view of γ, the amortized complexity is much better, as we shall show presently.

If l' lies on a fat edge (which is then the last edge of the left boundary) we proceed differently. This procedure is quite involved (see the full paper for details). Let us denote by B_γ the total time for crossing these fat edges when computing the view of γ. Obviously $B_\gamma = O(k \log n)$, if the view of γ consists of k tangencies and we *cross* the faces of the visible zone. In general we have:

Theorem 6 *Consider a convex object γ lying outside the convex hull of the set of obstacles, whose view consists of k tangencies. The total time needed to compute this view is*
(i) $O(k \log n)$, if we use the method of crossing faces;
(ii) $O(\log n + k + B_\gamma)$, if we traverse the visible zone of γ.

Proof (Sketch of ii) Consider the set F of faces intersected by the canonical image of γ. Let E be the set of edges passed during a traversal lying in the *left* boundary of a face in F. We shall prove $E = O(k)$. To this end consider edge $e \in E$ in the left boundary of $f_e \in F$. Let l be its maximal endpoint. The maximal free line segment corresponding to l is tangent to two objects, O_1 and O_2 say, at p_1 and p_2, respectively. Let l be directed from p_1 to p_2. If e is not the last edge we traverse in the left boundary of f_e, then it is not hard to see that p_2 is visible from γ along some ray r. We charge the cost of traversing e to the face f containing r. Note that $f \in F$. In this scheme every face of F is charged at most once, Since (i) vertex l is the *minimal* vertex of f, so it defines f uniquely, and (ii) l is the *maximal* endpoint of two edges, of which only one belongs to E. Also cf. [CG89] for a similar argument. $\square$

Remark 7 The standing hypothesis still is that convex objects have complexity $O(1)$. However, if the convex object γ consists of m arcs and line segments, we can prove that the view of γ can be determined in $O(\log n + k + m + B_\gamma)$ time, again provided γ lies outside the convex hull. This extension will be used in section 4.

4 Computing the Visibility Complex

Surgery on subdivisions

We present an example of a very simple situation, in which we explain the crucial step of the construction. The situation is defined mainly by pictures. Rigorous definitions are given later.

Consider three convex objects O_1, O_2 and O_3, such that O_3 lies outside the convex hull $CH(O_1 \cup O_2)$ of $O_1 \cup O_2$, see Figure 6a. The convex chain σ_0 is the boundary of the convex hull of $O_1 \cup O_2$, cut at its minimal point m_0. The subcomplex $\mathcal{P}(\sigma_0)$ is depicted in Figure 6b. There are two faces, whose points correspond to upward rays emanating from σ_0 along which we see either O_3 or 'the blue sky' O_∞.

The idea behind the construction of the visibility complex is to extend free space by adding pseudo–triangles. In our example we start with free space being the complement of $CH(O_1 \cup O_2) \cup O_3$. We then add the pseudo–triangle $a_0a_1a_2$, see Figure 6b. This amounts to removing the bitangent $s = a_0a_1$. The sides a_0a_2 and a_1a_2 then become part of the boundary of free space. Before removing s the subcomplex of chain a_0a_2 consists of two patches, see Figure 6b. (Here t_{ij} is the line supporting a_ia_j. Also compare with example 4.) Due to the removal of s the patch labeled s will change: along rays corresponding to points of this region we either see O_3 or the blue sky O_∞.

Upon removal of s chain σ_0 is split into two parts. The part with endpoint a_0 is concatenated to chain a_0a_2 to form a new chain m_0a_2, see Figure 6c. The subcomplex associated with this new chain is constructed from $\mathcal{P}(\sigma_0)$ and $\mathcal{P}(a_0a_2)$ by *surgery*. More precizely we cut $\mathcal{P}(\sigma_0)$ along the canonical image of chains a_0a_2 and a_2a_1 into two pieces. Piece $\mathcal{P}_1$ corresponds to the set of lines that intersect either (i) a_2a_0 and s (in this order), or (ii), or m_0a_0. Lines characterized by (i) form the shaded patch of $\mathcal{P}_1$. Considered as rays emanating from a_2a_0 they belong to the shaded patch of $\mathcal{P}(a_2a_0)$, see Figure 6b. Therefore the subcomplex of chain m_0a_2 is obtained by replacing the shaded patch of $\mathcal{P}(a_0a_2)$, labeled s in figure 6, with piece $\mathcal{P}_1$ of $\mathcal{P}(\sigma_0)$.

This example introduces many of the features of our method. First we introduce pseudo–triangulations of free space. The initial visibility complex is the collection of subcomplexes of the sides of the pseudo–triangle (cf. example 4). Processing a pseudo–triangle amounts to updating the visibility complex. We do this in two passes. In the first pass we update the subcomplex of the sides of the pseudo-triangle. This amounts to updating the view along

(a) A chain σ_0 and its subcomplex $\mathcal{P}(\sigma_0)$.

(b) Subcomplex of side $\sigma_1 = a_0a_2$ of pseudo-triangle $a_0a_1a_2$.

Figure 6: (c) Cutting the subcomplex $\mathcal{P}(\sigma_0)$ into patches $\mathcal{P}_1$ and $\mathcal{P}_2$ along the curve that is the canonical image of the sides a_0a_1 and a_1a_2 of pseudo–triangle $a_0a_1a_2$. The subcomplex of chain $m_0m_1a_2$ is obtained by replacing the patch labeled s (the shaded patch in part b) of $\mathcal{P}(\sigma_1)$ with patch $\mathcal{P}_1$ (the shaded patch in part c).

rays in free space that *leave* the pseudo–triangle, as in the example above. In the second pass we update the view along rays that *enter* the pseudo–triangle. We did not consider this part in the example above, but it involves e.g. updating the subcomplex of O_3. Starting with empty free space, we add these pseudo–triangles in a specific order, that allows for efficient update of the visibility complex. A triangulation that admits such an order is called *admissible*, a concept to be defined in the next subsection.

Admissible pseudo-triangulations

A *pseudo-triangulation* of the set of convex obstacles is a subdivision of the convex hull of the set of obstacles, such that every region is either the interior of an obstacle or a pseudo-triangle.

Our construction of the visibility complex starts with a special kind of pseudo–triangulation $\mathcal{T}$. For a non-horizontal line segment t let $\mathcal{T}_t$ be the sequence of pseudo–triangles intersected by t, ordered according to increasing y–coordinate.

Definition 8 *An admissible pseudo–triangulation is a pair $(\mathcal{T}, \lhd)$, consisting of a pseudo–triangulation $\mathcal{T}$ and a linear order $\lhd$ on the set of pseudo–triangles of $\mathcal{T}$ satisfying the following conditions.*
(i) For any free non–horizontal line segment t the sequence $\mathcal{T}_t$ is **unimodal** *with respect to $\lhd$: it is an increasing prefix followed by a decreasing suffix.*
(ii) If both endpoints of t are tangent to some obstacle, then $\mathcal{T}_t$ is **decreasing** *with respect to $\lhd$.*
(iii) Among the pseudo–triangles incident upon a convex obstacle O there are two pseudo–triangles R_0 and R_1 such that going along the boundary of O from R_0 to R_1 we pass a sequence of pseudo–triangles that is increasing with respect to $\lhd$, irrespective of whether we go clockwise or counterclockwise . (More precizely, R_0 (R_1) is the pseudo–triangle preceding the first **counterclockwise** *tangent we meet when walking in counterclockwise direction along ∂O, starting at the maximal (minimal) point of O.)*

Figure 7 shows an admissible pseudo-triangulation. It is intuitively clear that conditions (i) and (ii) are satisfied. The sequences referred to in condition (iii) associated with e.g. object O_1 are R_2, R_6 (clockwise) and R_2, R_3, R_4, R_5, R_6 (counterclockwise). Conditions (i) and (ii) give us control over the order in which the vertices of the visibility complex are computed. Condition (iii) is essential for achieving the optimal time bound. Without proof we state:

Proposition 9 *Any pseudo-triangulation of a scene of n convex obstacles in general position has $2n-2$ pseudo-triangles. There are $3n-3$ bitangents. There is an admissible pseudo-triangulation $(\mathcal{T}, \lhd)$, that can be constructed in $O(n \log n)$ time.*

Figure 7: An admissible pseudo–triangulation. The sequence $R_1, \cdots, R_6$ corresponds to the linear order $\lhd$.

The initial visibility complex

After the construction of an admissible pseudo–triangulation the algorithm sets up the subcomplexes associated with the chains forming the boundaries of the pseudo-triangles. These subcomplexes will be *augmented* by patches corresponding to rays emanating from a chain along which a bitangent in the boundary of the pseudo–triangle can be seen.

To be more precize consider a pseudo–triangle R with vertices a_0, a_1 and a_2. The subcomplex $\mathcal{P}(a_0a_1)$, associated with chain a_0a_1, is depicted in Figure 8b. Also compare Figure 5.

Let $s_0, \cdots, s_{h-1}$ and $s_h, \cdots, s_{k-1}$ be the counterclockwise sequences of bitangents contained in a_0a_2 and a_2a_1, respectively. Elements of these sequences will be called *facing left* and *facing right*, respectively. Let l_i be the directed line supporting s_i, a tangent line of either a_0a_2 or a_2a_1. The sequence $l_0, \cdots, l_{k-1}$ is ordered according to *decreasing* slope, so it corresponds to a decreasing sequence of points on the canonical image of a_0a_2 and a_2a_1, that can be determined in time $O(k)$.

If s_i is facing right (left) then u_i and v_i are the lines of smallest (largest) slope through the (counterclockwise successive) endpoints of s_i that intersect the chain a_0a_1. The sequence $u_0, v_0, \cdots, u_{k-1}, v_{k-1}$ is ordered according to *increasing* slope, with the possible exception of a prefix consisting of lines of slope 0 or a suffix consisting of lines of slope π. Therefore it corresponds to an increasing sequence of points on the canonical image of the chains a_0a_2 and a_2a_1, extended by the vertical segments $h_{a_0}h_m$ and $\overline{h}_m\overline{h}_{a_1}$,

Figure 8: The initial subcomplex of chain a_0a_1 of pseudo–triangle $a_0a_1a_2$.

see Figure 8. This sequence can also be determined in $O(k)$ time.

Patches $P(s_0), \cdots, P(s_{h-1})$ of the subcomplex of a_0a_1 correspond to rays along which we can see the *right* facing bitangents $s_0, \cdots, s_{h-1}$. They are subsets of the leftmost region of $\mathcal{P}(a_0a_1)$, and are constructed by introducing a sequence of pairwise disjoint edges $u_0l_0, v_0l_0, \cdots, u_{h-1}l_{h-1}, v_{h-1}l_{h-1}$ in this leftmost region of $\mathcal{P}(a_0a_1)$. Patch $P(s_i)$, $0 \leq i < h$, is the region bounded by edges u_il_i and v_il_i and the boundary of $\mathcal{P}(a_0a_1)$. Note that edge u_0l_0 may coincide with a boundary edge of $\mathcal{P}(a_0a_1)$ if s_0 is tangent to a_0a_1, in which case $u_0 = h_{a_0}$ and $l_0 = t_0$. Similarly edge $v_{h-1}l_{h-1}$ may be degenerate: if s_{h-1} is tangent to a_0a_2 we have $v_{h-1} = l_{h-1} = t_2$. In this case $P(s_{h-1})$ is a 'digonal' patch.

Patches $P(s_h), \cdots, P(s_{k-1})$, corresponding to the *left* facing bitangents $s_h, \cdots, s_{k-1}$, lie in the rightmost region of $\mathcal{P}(a_0a_1)$ and are constructed similarly.

Since the sequence of edges bounding these patches is ordered we can construct these patches in $O(k)$ time. There are $O(n)$ bitangents in the pseudo-triangulation, each giving rise to $O(1)$ patches. Therefore the overall time needed for the construction of the subcomplexes of all chains in the pseudo–triangulation is $O(n)$. Hence we have proved:

Lemma 10 *Given an admissible pseudo–triangulation of a set of n convex obstacles, the initial visibility complex, consisting of the collection of augmented subcomplexes of all chains in the pseudo–triangulation, can be constructed in $O(n)$ time.*

Pass 1

As announced in the previous section we process the pseudo–triangles in the admissible order. Let the current pseudo–triangle be $R = a_0a_1a_2$, and let $s_0, \cdots, s_{k-1}$ again be the counterclockwise sequence of bitangents contained in a_0a_2 and a_2a_1, see Figure 8. Let R_i be the pseudo–triangle sharing s_i with R. In this first pass we show how to 'fill in' patches $P(s_i)$ of $\mathcal{P}(a_0a_1)$, $\mathcal{P}(a_1a_2)$ and $\mathcal{P}(a_2a_0)$ for values of i such that R_i precedes R.

Consider patch $P(s_i)$ of $\mathcal{P}(a_0a_1)$ corresponding to such a bitangent s_i, see Figure 8. First consider the case in which s_i is *not tangent* to a_0a_1 (so $i \neq 0$ in the situation of Figure 8). The subcomplex $\mathcal{P}_i$ of the chain containing s_i is depicted in Figure 9. Here p_i and q_i are the endpoints of s_i. (In our example s_i is a chain by itself, but it might be a proper subset of a chain if the tangent at a_2 contains s_i.) As in the

Figure 9: The subcomplex $\mathcal{P}_i$ and the canonical image γ_i of a_0a_1 and a_1a_2. Region $l_iu_iv_i$ (shaded) replaces region $P(s_i) = l_iu_iv_i$ of $\mathcal{P}(a_0a_1)$, see Figure 8. Region $h_{p_i}u_iv_ih_{q_i}$ replaces a similar region in $\mathcal{P}(a_1a_2)$.

example of section 4 we do surgery on $\mathcal{P}_i$ by cutting it along the canonical image γ_i of a_0a_1 and a_2a_1. The cut is bounded by u_i and v_i, and passes several faces of the subcomplex, see Figure 9. The points of intersection of γ_i and edges of $\mathcal{P}_i$ are 'new' vertices of the visibility complex (arcs of the tangent visibility graph).

The crucial observation is that all vertices of the visibility complex, corresponding to an upwardly directed free bitangent (and hence an arc of the tangent visibility graph) emanating from a side of the current pseudo–triangle R, are detected during pass 1 when processing R. This is an obvious consequence of the fact that the pseudo–triangulation we start with is *admissible*.

Note that γ_i represents a segment of the view of a_0a_1 and a_2a_1, bounded by rays u_i and v_i. Therefore the intersection of γ_i and the subdivision $\mathcal{P}_i$ is determined in exactly the same way as we found the view of an object *outside the convex hull*, see section 3, provided we can efficiently determine the initial intersection, viz. the first intersection when starting from u_i. To this end we store the sequences of edges of the pair of 'outer' faces, viz. the faces incident upon the boundary of $\mathcal{P}_i$, in a concatenable queue. This enables us to find the first intersection in $O(\log n)$ time. Moreover these data structures can be maintained after surgery in $O(\log n)$ time as well. It is easy to see that this $O(\log n)$ cost is paid $O(1)$ time per bitangent s_i of the pseudo–triangulation, adding up to $O(n \log n)$ time overall. Note that a_0a_1 and a_1a_2 do not necessarily have complexity $O(1)$. However, in view of theorem 6 and remark 7, the time needed to find the other intersections is proportional to the number of new vertices lying on γ_i, plus the number of arcs in a_0a_1 and a_2a_1, plus the time needed to cross 'fat' edges. The total number of arcs in the initial pseudo–triangulation is $O(n)$, however, so this contribution does not dominate the total time complexity.

Finally consider the situation in which s_i is tangent to a_0a_1, say at a_0 (so $i = 0$). Let s_0 belong to a chain σ_0 that extends beyond a_0. As in the example of section 4 this chain is split by removing s_0, and the part that ends at a_0 is concatenated to a_1a_0 to form a new chain, $\overline{\sigma_0}$ say. We find $\mathcal{P}(\overline{\sigma_0})$ again by cutting $\mathcal{P}(\sigma_0)$ along the canonical image γ_0 of a_0a_1 and a_2a_1. (In the situation of figure 8 γ_0 is bounded by h_m and v_0.) Subsequently we replace the region $P(s_0)$ of $\mathcal{P}(a_0a_1)$ with one of the pieces of $\mathcal{P}(\sigma_0)$. Note that processing s_0 is completely similar to processing s in the example of section 4.

Summarizing the previous discussion we have proved

Lemma 11 *The total time needed to perform pass 1 on all pseudo–triangles is $O(n \log n + k + B)$, where k is the number of arcs of the tangent visibility graph and B is the time needed to cross the 'fat' edges.*

Although at first glance we can't beat $B = O(k \log n)$, we shall sketch in section ?? how to amortize the total cost of traversing the 'fat' edges in such a way that $B = O(n)$.

Pass 2

Let R again be the current pseudo–triangle, sharing a bitangent $s = pq$ with a pseudo–triangle R' which has already been processed, and therefore precedes it in the admissible order. Let σ_0 be the chain in ∂R containing pq. The lower endpoint of s is p, the upper endpoint is q. We assume that s is facing right, see Figure 10. Let σ_1 be the chain in the boundary of the region on the other side of s from which s can be seen along horizontal rays arbitrarily near p. Due to the removal of s the subcomplex $\mathcal{P}(\sigma_1)$ needs to be updated. As in pass this is done by surgery on the subcomplex $\mathcal{P}(\sigma_0)$. However, in the present situation the surgery is simpler, since there are no upwardly directed free bitangents that intersect s from left to right. In other words: we don't find new vertices in this pass, we merely move them from one subcomplex, $\mathcal{P}(\sigma_0)$ in this case, to an other one, $\mathcal{P}(\sigma_1)$.

Figure 10: Pass 2 of processing triangle R. We distinguish whether R and the obstacle containing p lie on the same side or on different sides of the line t supporting $s = pq$.

Lemma 12 *The canonical image of p in $\mathcal{P}(\sigma_1)$ is a curve, bounded by points corresponding to t and $\overline{h}_p$, that lies in a single face of $\mathcal{P}(\sigma_1)$.*

Proof. Suppose there is an edge of $\mathcal{P}(\sigma_1)$ that intersects the canonical image of p in a point corresponding to a line l through p. Obviously p is not an endpoint of σ_1, so we are in the situation depicted in the left part of figure 10. But then a slight perturbation of l yields a free line segment l' that is tangent to two objects on different sides of s, but intersects R' *before* R. This contradicts the fact that the pseudo–

triangulation is admissible, cf. definition 8(ii), which proves the lemma. □

Let γ be the curve in $\mathcal{P}(\sigma_1)$ that is the lower boundary of the face containing the canonical image of p. Let m be the extremal point that is hit first when moving a horizontal ray connecting σ_0 and s in upward direction from p. If no such point exists we take $m = q$. Then γ is a curve connecting t with $\overline{h}_m$. Let $\overline{(\gamma)}$ be the corresponding curve in $\mathcal{P}(\sigma_0)$: its points are rays emanating from s that arise when we extend rays corresponding to points of γ.

Lemma 13 *(i) The line t corresponds to a point in the boundary of $\mathcal{P}(\sigma_0)$.*
(ii) The curve $\overline{\gamma}$ lies in a single face of $\mathcal{P}(\sigma_0)$.

The proof is similar to that of lemma 12. It is not hard to see that the piece of $\mathcal{P}(\sigma_1)$ between γ and the canonical image of p can be updated by surgery: cut $\mathcal{P}(\sigma_0)$ along $\overline{\gamma}$, and insert it into $\mathcal{P}(\sigma_1)$ by identifying points of the cut $\overline{\gamma}$ with the corresponding points of γ. Lemma 13(i) guarantees that this surgery yields two pieces, while (ii) shows that we can perform this surgery in $O(\log n)$ time by splitting the boundary of the face containing $\overline{\gamma}$ at t. If σ_2 is the chain from which q can be seen along a horizontal ray with slope π we update $\mathcal{P}(\sigma_2)$ near the canonical image of q in a completely similar fashion. There are $O(n)$ bitangents in an admissible pseudo–triangulation, so:

Lemma 14 *The total time needed to execute pass 2 for all pseudo–triangles is $O(n \log n)$.*

Complexity of the algorithm

In view of proposition 9 and lemma's 10, 11 and 14 we see that the total time complexity of our construction is $O(n \log n + k)$, provided we can prove that the total time B needed to cross the 'fat' edges is $O(n)$. Due to lack of space we merely mention that a completely similar problem, although in a different disguise, occurs in [GM91]. Here a clever use of a *split–find* data structure invented by Gabow and Tarjan yields an amortization scheme with $B = O(n)$. In this version we merely mention that the technical condition (iii) in definition 8 is crucial for this approach. Summarizing we have proved:

Theorem 15 *The visibility complex $\mathcal{P}$ of a collection of n pairwise disjoint convex obstacles can be constructed in $O(n \log n + k)$ time, where k is the size of $\mathcal{P}$ (or, equivalently, of the tangent visibility graph).*

5 Conclusion

We expect that our methods can be used to solve various other geometric problems, like e.g. planning the motion of a rod amidst convex obstacles, cf. [Veg91], ray shooting (this will require a persistent data structure for the visibility complex, cf. [Poc90], and the computation of a sector of the visibility polygon. Another interesting question is concerned with classification: although some partial results are known a complete classification of visibility graphs still seems to be lacking. Due to the richer structure it might give more insight into the problem of classifying cell complexes that are visibility complexes.

References

[CG89] B. Chazelle and L. Guibas. Visibility and intersection problems in plane geometry. *Discrete and Computational Geoemetry*, 4:551–581, 1989.

[GM91] S. K. Ghosh and D. Mount. An output sensitive algorithm for computing visibility graphs. *Siam J. Comput.*, 20:888–910, 1991.

[Poc90] M. Pocchiola. Graphics in flatland revisited. In *Proceedings of the 2nd Scandinavian Workshop on Algorithm Theory, Bergen*, pages 85–96. Springer Verlag, Juillet 1990.

[Veg90] G. Vegter. The visibility diagram: a data structure for visibility and motion planning problems in the plane. In *Proceedings of the 2nd Scandinavian Workshop on Algorithm Theory, Bergen*, pages 97–110. Springer Verlag, Juillet 1990.

[Veg91] G. Vegter. Dynamically maintaining the visibility graph. In *Proceedings of the 2nd Workshop WADS'91, Ottawa, Canada*, pages 425–436. Springer Verlag, August 1991.

[Wel85] E. Welzl. Constructing the visibility graph of n lines segments in the plane. *Information Processing Letters*, 20:167–171, 1985.